

# HP StorageWorks

## X9300 Network Storage Gateway Administrator Guide

### Abstract

This guide describes tasks related to cluster configuration and monitoring, system upgrade and recovery, hardware component replacement, and troubleshooting. It does not document X9000 file system features or standard Linux administrative tools and commands. For information about configuring and using X9000 Software file system features, see the *HP StorageWorks X9000 File Serving Software File System User Guide*.

This guide is intended for system administrators and technicians who are experienced with installing and administering networks, and with performing Linux operating and administrative tasks.



## Legal and notice information

© Copyright 2010 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

## Acknowledgments

Microsoft, Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

## Warranty

WARRANTY STATEMENT: To obtain a copy of the warranty for this product, see the warranty information website:

<http://www.hp.com/go/storagewarranty>

## Revision History

Edition	Date	Software Version	Description
First	April 2010	5.4	Initial release of the X9300 Gateway.
Second	August 2010	5.4.1	Added management console backup, migration to an agile management console configuration, software upgrade procedures, and system recovery procedures.
Third	August 2010	5.4.1	Revised the upgrade procedure.
Fourth	December 2010	5.5	Added information about NDMP backups and virtual interface configuration, and updated cluster procedures.
Fifth	March 2011	5.5	Updated segment evacuation procedure.

---

# Contents

<b>1 Product description .....</b>	<b>9</b>
HP X9300 Network Storage Gateway features .....	9
System components .....	9
HP X9000 Software features .....	10
High availability and redundancy .....	10
<b>2 Getting started .....</b>	<b>13</b>
Setting up the X9300 Network Storage Gateway .....	13
Installation steps .....	13
Additional configuration steps .....	13
Management interfaces .....	14
Using the GUI .....	14
Customizing the GUI .....	17
Adding user accounts for GUI access .....	18
Using the CLI .....	18
Starting the array management software .....	18
X9000 client interfaces .....	18
X9000 Software manpages .....	19
Changing passwords .....	19
<b>3 Configuring virtual interfaces for client access .....</b>	<b>21</b>
Network and VIF guidelines .....	21
Creating a bonded VIF .....	21
Configuring standby backup nodes .....	22
Configuring NIC failover .....	22
Configuring automated failover .....	22
Example configuration .....	22
Specifying VIFs in the client configuration .....	23
<b>4 Configuring failover .....</b>	<b>25</b>
Agile management consoles .....	25
Agile management console modes .....	25
Agile management consoles and failover .....	25
Viewing information about management consoles .....	26
Cluster high availability .....	26
Failover modes .....	26
What happens during a failover .....	27
Setting up automated failover .....	27
Identifying standbys for file serving nodes .....	27
Identifying power sources .....	28
Turning automated failover on and off .....	30
Manually failing over a file serving node .....	30
Failing back a file serving node .....	30
Using network interface monitoring .....	31

Setting up HBA monitoring .....	33
Discovering HBAs .....	33
Identifying standby-paired HBA ports .....	34
Turning HBA monitoring on or off .....	34
Deleting standby port pairings .....	34
Deleting HBAs from the configuration database .....	34
Displaying HBA information .....	34
Checking the High Availability configuration .....	35
<b>5 Configuring cluster event notification .....</b>	<b>37</b>
Setting up email notification of cluster events .....	37
Associating events and email addresses .....	37
Configuring email notification settings .....	37
Turning email notifications on or off .....	38
Dissociating events and email addresses .....	38
Testing email addresses .....	38
Viewing email notification settings .....	38
Setting up SNMP notifications .....	38
Configuring the SNMP agent .....	39
Configuring trapsink settings .....	40
Associating events and trapsinks .....	40
Defining views .....	41
Configuring groups and users .....	41
Deleting elements of the SNMP configuration .....	42
Listing SNMP configuration information .....	42
<b>6 Configuring system backups .....</b>	<b>43</b>
Backing up the management console configuration .....	43
Using NDMP backup applications .....	43
Configuring NDMP parameters on the cluster .....	44
NDMP process management .....	45
Viewing or canceling NDMP sessions .....	45
Viewing or rescanning tape and media changer devices .....	46
NDMP events .....	46
<b>7 Creating hostgroups for X9000 clients .....</b>	<b>47</b>
How hostgroups work .....	47
Creating a hostgroup tree .....	47
Adding an X9000 client to a hostgroup .....	48
Adding a domain rule to a hostgroup .....	48
Viewing hostgroups .....	49
Deleting hostgroups .....	49
Other hostgroup operations .....	49
<b>8 Monitoring cluster operations .....</b>	<b>51</b>
Monitoring the status of file serving nodes .....	51
Monitoring cluster events .....	51
Viewing events .....	52
Removing events from the events database table .....	52
Monitoring cluster health .....	52
Health checks .....	52
Health check reports .....	53
Viewing logs .....	55

Viewing operating statistics for file serving nodes .....	56
<b>9 Maintaining the system .....</b>	<b>57</b>
Shutting down the system .....	57
Shutting down the X9000 Software .....	57
Powering off the hardware .....	58
Starting the system .....	58
Starting the X9000 Software .....	58
Powering file serving nodes on or off .....	58
Starting and stopping processes .....	59
Tuning file serving nodes and X9000 clients .....	59
Migrating segments .....	60
Removing storage from the cluster .....	61
Maintaining networks .....	63
Cluster and user network interfaces .....	63
Adding user network interfaces .....	63
Setting network interface options in the configuration database .....	64
Preferring network interfaces .....	65
Unpreferring network interfaces .....	65
Making network changes .....	66
Changing the IP address for a Linux X9000 client .....	66
Changing the IP address for the cluster interface on a dedicated management console .....	66
Changing the cluster interface .....	66
Managing routing table entries .....	67
Deleting a network interface .....	67
Viewing network interface information .....	67
<b>10 Migrating to an agile management console configuration .....</b>	<b>69</b>
Backing up the configuration .....	69
Performing the migration .....	69
Removing the dedicated Management Server .....	72
<b>11 Upgrading the X9000 Software .....</b>	<b>75</b>
Automatic upgrades .....	75
Manual upgrades .....	76
Standard upgrade for clusters with a dedicated Management Server machine or blade .....	77
Standard online upgrade .....	77
Standard offline upgrade .....	79
Agile upgrade for clusters with an agile management console configuration .....	81
Agile online upgrade .....	81
Agile offline upgrade .....	85
Upgrading Linux X9000 clients .....	88
Upgrading Windows X9000 clients .....	89
Troubleshooting upgrade issues .....	89
<b>12 Licensing .....</b>	<b>91</b>
Viewing license terms .....	91
Retrieving a license key .....	91
Using AutoPass to retrieve and install permanent license keys .....	91
<b>13 Upgrading firmware .....</b>	<b>93</b>
Upgradable firmware .....	93

Installing firmware upgrades .....	93
<b>14 Troubleshooting .....</b>	<b>95</b>
Managing support tickets .....	95
Creating, viewing, and deleting support tickets .....	95
Support ticket states .....	96
Updating the ticket database when nodes are added or removed .....	96
Configuring the support ticket feature .....	96
Configuring shared ssh keys .....	97
Viewing software version numbers .....	97
Troubleshooting specific issues .....	98
Software services .....	98
Failover .....	98
Windows X9000 clients .....	99
Synchronizing information on file serving nodes and the configuration database .....	99
<b>15 Replacing components .....</b>	<b>101</b>
Customer replaceable components .....	101
Hot-pluggable and non-hot-pluggable components .....	101
Returning the defective component .....	101
Parts-only warranty service .....	102
Required tools .....	102
Additional documentation .....	102
Replacing a system board .....	102
Replacing a NIC adapter .....	103
Replacing a Fibre Channel HBA .....	104
<b>16 Recovering a file serving node .....</b>	<b>107</b>
Starting the recovery .....	107
Configuring a file serving node using the original template .....	108
Completing the restore on a file serving node .....	112
Configuring a file serving node manually .....	112
<b>17 Support and other resources .....</b>	<b>123</b>
Contacting HP .....	123
Related information .....	123
HP websites .....	124
Rack stability .....	124
Customer self repair .....	124
<b>A Component and cabling diagrams .....</b>	<b>125</b>
Component diagrams .....	125
Front view of file serving node .....	125
Rear view of file serving node .....	125
Cabling diagrams .....	128
Cluster network cabling diagram .....	128
<b>B Spare parts list .....</b>	<b>129</b>
1GbE spare parts .....	129
1 GbE (AW539A) .....	129
1 GbE (AW539B) .....	130
10 GbE spare parts .....	132

10 GbE (AW540A) .....	132
10 GbE/IB (AW540B) .....	133
IB (AW541A) .....	134
Base rack (AW546A) .....	135

## C Warnings and precautions ..... 137

Electrostatic discharge information .....	137
Grounding methods .....	137
Equipment symbols .....	137
Rack warnings and precautions .....	138
Device warnings and precautions .....	140

## D Regulatory compliance and safety ..... 143

Regulatory compliance identification numbers .....	143
Federal Communications Commission notice .....	143
Class A equipment .....	143
Class B equipment .....	143
Declaration of conformity for products marked with the FCC logo, United States only .....	144
Modifications .....	144
Cables .....	144
Laser compliance .....	144
International notices and statements .....	145
Canadian notice (Avis Canadien) .....	145
Class A equipment .....	145
Class B equipment .....	145
European Union notice .....	145
BSMI notice .....	146
Japanese notice .....	146
Korean notice (A&B) .....	146
Safety .....	147
Battery Replacement notice .....	147
Taiwan Battery Recycling Notice .....	147
Power cords .....	147
Japanese Power Cord notice .....	148
Electrostatic discharge .....	148
Preventing electrostatic discharge .....	148
Grounding methods .....	148
Waste Electrical and Electronic Equipment directive .....	149
Czechoslovakian notice .....	149
Danish notice .....	149
Dutch notice .....	149
English notice .....	150
Estonian notice .....	150
Finnish notice .....	150
French notice .....	151
German notice .....	151
Greek notice .....	151
Hungarian notice .....	152
Italian notice .....	152
Latvian notice .....	152
Lithuanian notice .....	153
Polish notice .....	153
Portuguese notice .....	153
Slovakian notice .....	154

Slovenian notice .....	154
Spanish notice .....	154
Swedish notice .....	155

Glossary .....	157
----------------	-----

Index .....	161
-------------	-----



---

# 1 Product description

The HP X9300 Network Storage Gateway is a flexible, scale-out solution that brings feature-rich gateway file services to HP MSA, EVA, P4000, or 3rd-party arrays or SANs.



## HP X9300 Network Storage Gateway features

The X9300 Network Storage Gateway provides the following features:

- Segmented, scalable file system under a single namespace
- NFS, CIFS, FTP, and HTTP support for accessing file system data
- Centralized CLI and GUI cluster management
- Policy management
- Continuous remote replication

---

### ❗ IMPORTANT:

It is important to keep regular backups of the cluster configuration.

---

## System components

System components include:

- Optional HP StorageWorks X9300 Network Storage System Base Rack:
  - Keyboard, video, and mouse (KVM)
  - HP ProCurve 2910-24G management switch
- HP StorageWorks X9300 file serving node
- Storage connectivity (one or more options chosen by the customer):
  - Dual port 10-GbE NIC (iSCSI)
  - Quad port 1-GbE NIC (iSCSI)
  - Dual port FC 8-GB HBA (Fibre Channel)
  - SAS HBA
- Storage array support:
  - iSCSI—HP LeftHand P4000, Dell EqualLogic
  - Fibre Channel—HP EVA, HP MSA2000, EMC CLARiiON
  - SAS—HP MSA

- Front-end connectivity (chosen by the customer):
  - 10-GbE network
  - 1-GbE network
  - Infiniband network
- Pre-installed software:
  - Red Hat Linux operating system
  - HP StorageWorks X9000 File Serving Software
  - Integrated Lights-Out 2 (iLO 2) remote management software

For component and cabling diagrams, see [Appendix A](#).

## HP X9000 Software features

HP X9000 Software is a scale-out, network-attached storage solution composed of a parallel file system for clusters, an integrated volume manager, high-availability features such as automatic failover of multiple components, and a centralized management interface. X9000 Software can be deployed in environments scaling to thousands of nodes.

Based on a Segmented File System architecture, X9000 Software enables enterprises to integrate I/O and storage systems into a single clustered environment that can be shared across multiple applications and managed from a single central management console.

X9000 Software is designed to operate with high-performance computing applications that require high I/O bandwidth, high IOPS throughput, and scalable configurations. Examples of these applications include Internet streaming, rich media streaming, data mining, web search, manufacturing, financial modeling, life sciences modeling, and seismic processing.

Some of the key features and benefits are as follows:

- Scalable configuration. You can add servers to scale performance and add storage devices to scale capacity.
- Single namespace. All directories and files are contained in the same namespace.
- Multiple environments. Operates in both the SAN and DAS environments.
- High availability. The high-availability software protects servers.
- Tuning capability. The system can be tuned for large or small-block I/O.
- Flexible configuration. Segments can be migrated dynamically for rebalancing and data tiering.

## High availability and redundancy

The segmented architecture is the basis for fault resilience—loss of access to one or more segments does not render the entire file system inaccessible. Individual segments can be taken offline temporarily for maintenance operations and then returned to the file system.

To ensure continuous data access, X9000 Software provides manual and automated failover protection at various points:

- **Server.** A failed node is powered down and a designated standby server assumes all of its segment management duties.
- **Segment.** Ownership of each segment on a failed node is transferred to a designated standby server.
- **Network interface.** The IP address of a failed network interface is transferred to a standby network interface until the original network interface is operational again.

- **Storage connection.** For servers with HBA-protected Fibre Channel access, failure of the HBA triggers failover of the node to a designated standby server.



---

## 2 Getting started

---

### ❗ IMPORTANT:

Do not modify any parameters of the operating system or kernel, or update any part of the X9300 Network Storage Gateway unless instructed to do so by HP; otherwise, the X9300 Network Storage Gateway could fail to operate properly.

---

## Setting up the X9300 Network Storage Gateway

### Installation steps

The HP StorageWorks X9300 Network Storage Gateway is deployed at your site according to the terms in your Statement of Work. See your Statement of Work for details.

### Additional configuration steps

When your system is up and running, you can perform any additional configuration of your cluster and file systems. The management console GUI and CLI are used to perform most operations. (Some of the features described here might have been configured for you as part of the system installation.)

**Cluster.** Configure the following as needed:

- Virtual interfaces for client access.
- Failover for file serving nodes, network interfaces, and HBAs.
- Cluster event notification through email or SNMP.
- Management console backups.
- NDMP backups.

These cluster features are described later in this guide.

**File systems.** Set up the following features as needed:

- Additional file systems. Optionally, configure data tiering on the file systems to move files to specific tiers based on file attributes.
- NFS, CIFS, FTP, or HTTP. Configure the methods you will use to access file system data.
- Quotas. Configure user, group, and directory tree quotas as needed.
- Remote replication. Use this feature to replicate changes in a source file system on one cluster to a target file system on either the same cluster or a second cluster.
- Snapshots. Use this feature to capture a point-in-time copy of a file system.
- File allocation. Use this feature to specify the manner in which segments are selected for storing new files and directories.

For more information about these file system features, see the *HP StorageWorks File Serving Software File System User Guide*.

# Management interfaces

Cluster operations are managed through the X9000 Software management console, which provides both a GUI and a CLI. Most operations can be performed from either the GUI or the CLI. However, the following operations can be performed only from the CLI:

- SNMP configuration (`ibrix_snmpagent`, `ibrix_snmpgroup`, `ibrix_snmptrap`, `ibrix_snmpuser`, `ibrix_snmpview`)
- Health checks (`ibrix_haconfig`, `ibrix_health`, `ibrix_healthconfig`)
- Raw storage management (`ibrix_pv`, `ibrix_vg`, `ibrix_lv`)
- Management console operations (`ibrix_fm`) and management console tuning (`ibrix_fm_tune`)
- File system checks (`ibrix_fsck`)
- Kernel profiling (`ibrix_profile`)
- NFS autoconnection (`ibrix_autoconnect`)
- Cluster configuration (`ibrix_clusterconfig`)
- Configuration database consistency (`ibrix_dbck`)
- Shell task management (`ibrix_shell`)

## Using the GUI

The GUI is a browser-based interface to the management console. See the release notes for the supported browsers and other software required to view charts on the dashboard.

If you are using HTTP to access the GUI, navigate to the following location, specifying port 80:

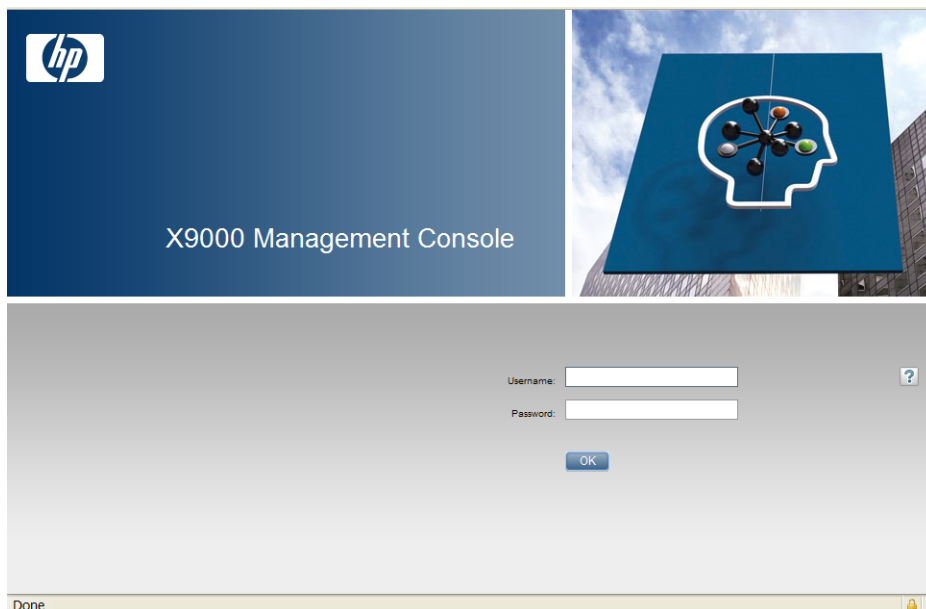
```
http://<management_console_IP>:80/fusion
```

If you are using HTTPS to access the GUI, navigate to the following location, specifying port 443:

```
https://<management_console_IP>:443/fusion
```

In these URLs, `<management_console_IP>` is the IP address of the management console user VIF.

The GUI prompts for your user name and password. The default administrative user is **ibrix**. Enter the password that was assigned to this user when the system was installed. (You can change the password using the Linux `passwd` command.) To allow other users to access the GUI, see [“Adding user accounts for GUI access”](#) on page 18.






The GUI dashboard opens in the same browser window. You can open multiple GUI windows as necessary. See the online help for information about all GUI displays and operations.



The GUI dashboard enables you to monitor the entire cluster. There are three parts to the dashboard: System Status, Cluster Overview, and the Navigator.

## System Status

The System Status section lists the number of cluster events that have occurred in the last 24 hours. There are three types of events:

	<b>Alerts.</b> Disruptive events that can result in loss of access to file system data. Examples are a segment that is unavailable or a server that cannot be accessed.
	<b>Warnings.</b> Potentially disruptive conditions where file system access is not lost, but if the situation is not addressed, it can escalate to an alert condition. Examples are a very high server CPU utilization level or a quota limit close to the maximum.
	<b>Information.</b> Normal events that change the cluster. Examples are mounting a file system or creating a segment.

## Cluster Overview

The Cluster Overview provides the following information:

### Capacity

The amount of cluster storage space that is currently free or in use.

### Filesystems



The current health status of the file systems in the cluster. The overview reports the number of file systems in each state (healthy, experiencing a warning, experiencing an alert, or unknown).

### Segment Servers

The current health status of the file serving nodes in the cluster. The overview reports the number of nodes in each state (healthy, experiencing a warning, experiencing an alert, or unknown).

### Services

Whether the specified file system services are currently running:

	One or more tasks are running.
	No tasks are running.

### Statistics

Historical performance graphs for the following items:

- Network I/O (MB/s)
- Disk I/O (MB/s)
- CPU usage (%)
- Memory usage (%)

On each graph, the X-axis represents time and the Y-axis represents performance.

Use the **Statistics** menu to select the servers to monitor (up to two), to change the maximum value for the Y-axis, and to show or hide resource usage distribution for CPU and memory.

### Recent Events

The most recent cluster events. Use the **Recent Events** menu to select the type of events to display.

You can also access certain menu items directly from the Cluster Overview. Mouse over the Capacity, Filesystems or Segment Server indicators to see the available options.

## Navigator

The Navigator appears on the left side of the window and displays the cluster hierarchy. You can use the Navigator to drill down in the cluster configuration to add, view, or change cluster objects such as file systems or storage, and to initiate or view tasks such as snapshots or replication. When



you select an object, a details page shows a summary for that object. The lower Navigator allows you to view details for the selected object, or to initiate a task. In the following example, we selected Cluster Configuration in the Navigator, and the Summary shows configuration information. In the lower Navigator, we selected **NDMP Backup > Active Sessions** to see details about the sessions.

Cluster:

User: root Role: admin Options

System Status

Updated Jun. 28, 2010, 12:35:45 PM MDT

Event Status (24 hours): 4 1 5788

Navigator

- Dashboard
- Filesystems
- Servers
- Storage
- Vendor Storage
- Events
- Clients
- FTP
- Certificates
- Hostgroups
- Cluster Configuration
- HTTP
- Support Tickets

Summary

Name	Value
Total Filesystem Count	1
Total Filesystems in Alert State	0
Total Filesystems in Warning State	0
Total Filesystems in OK State	1
Total Filesystem Blocks (GB)	9.05
Free Filesystem Blocks (GB)	8.95
Used Filesystem Blocks (GB)	0.11
Used Filesystem Blocks (%)	1.18
Free Filesystem Blocks (%)	98.82
Total Servers Count	2
Total Servers in Alert State	0
Total Servers in Warning State	0
Total Servers in OK State	2
Servers Network Throughput (MB/s)	0.00
Servers I/O Throughput (MB/s)	0.00
Alert Event Count	4
Warning Event Count	1
Info Event Count	5788

Cluster

- Cluster
- Email
- SNMP
- NDMP Backup
- Active Sessions
- Session History
- Tape Devices
- License

Active NDMP Sessions

Options

Hostname	Identifier	Type	Start Time	DMA IP Address
lrmv2	123	FOO	Wed May 26 22:46:39 2010	192.168.10.1
lrmv2	15543	IDLE	Wed May 26 22:46:39 2010	192.168.10.2
lrmv2	17371	IDLE	Wed May 26 22:46:39 2010	192.168.10.2
lrmv2	18129	IDLE	Wed May 26 22:46:39 2010	192.168.10.2
lrmv3	15543	IDLE	Wed May 26 22:46:39 2010	192.168.10.2

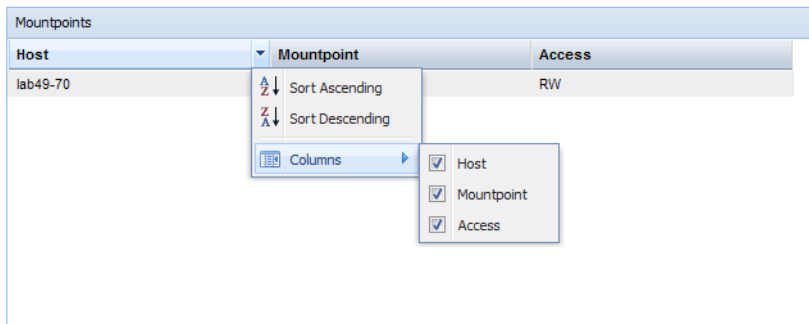


#### NOTE:

When you perform an operation on the GUI, a spinning finger is displayed until the operation is complete. However, if you use Windows Remote Desktop to access the management console, the spinning finger is not displayed.

## Customizing the GUI

For most tables in the GUI, you can specify the columns that you want to display and the sort order of each column. When this feature is available, mousing over a column causes the label to change color and a pointer to appear. Click the pointer to see the available options. In the following example, you can sort the contents of the Mountpoint column in ascending or descending order, and you can select the columns that you want to appear in the display.



## Adding user accounts for GUI access

X9000 Software supports administrative and user roles. When users log in under the administrative role, they can configure the cluster and initiate operations such as remote replication or snapshots. When users log in under the user role, they can view the cluster configuration and status, but cannot make configuration changes or initiate operations. The default administrative user name is `ibrix`. The default regular username is `ibrixuser`.

Username for the administrative and user roles are defined in the `/etc/group` file. Administrative users are specified in the `ibrix-admin` group, and regular users are specified in the `ibrix-user` group. These groups are created when X9000 Software is installed. The following entries in the `/etc/group` file show the default users in these groups:

```
ibrix-admin:x:501:root,ibrix
ibrix-user:x:502:ibrix,ibrixUser,ibrixuser
```

You can add other users to these groups as needed, using Linux procedures.

## Using the CLI

The administrative commands described in this guide must be executed on the management console host and require root privileges. The commands are located in `$IBRIXHOME/bin`. For complete information about the commands, see the *HP StorageWorks X9000 File Serving Software CLI Reference Guide*.

When using `ssh` to access the machine hosting the management console, specify the IP address of the management console user VIF.

## Starting the array management software

Depending on the array type, you can launch the array management software from the management console GUI. In the Navigator, select **Vendor Storage**, select your array from the Vendor Storage page, and click **Launch Storage Management**.

## X9000 client interfaces

X9000 clients can access the management console as follows:

- **Linux clients.** Linux client commands can be used for tasks such as mounting or unmounting file systems and displaying statistics. See the *HP StorageWorks X9000 File Serving Software CLI Reference Guide* for details about these commands.
- **Windows clients.** The Windows client GUI can be used for tasks such as mounting or unmounting file systems and registering Windows clients.

## Using the Windows X9000 client GUI

The Windows X9000 client GUI is the client interface to the management console. To open the GUI, double-click the desktop icon or select the IBRIX Client program from the Start menu on the client. The client program contains tabs organized by function.



### NOTE:

The Windows X9000 client application can be started only by users with Administrative privileges.

- **Status.** Shows the client's management console registration status and mounted file systems, and provides access to the IAD log for troubleshooting.
- **Registration.** Registers the client with the management console, as described in the *HP StorageWorks File Serving Software Installation Guide*.
- **Mount.** Mounts a file system. Select the Cluster Name from the list (the cluster name is the management console name), enter the name of the file system to mount, select a drive, and then click **Mount**. (If you are using Remote Desktop to access the client and the drive letter does not appear, log out and log back in.)
- **Unmount.** Unmounts a file system.
- **Tune Host.** Tunable parameters include the NIC to prefer (the client uses the cluster interface by default unless a different network interface is preferred for it), the communications protocol (UDP or TCP), and the number of server threads to use.
- **Active Directory Settings.** Displays current Active Directory settings.

Online help is also available for the client GUI.

## X9000 Software manpages

X9000 Software provides manpages for most of its commands. To view the manpages, set the `MANPATH` variable on the management console to include the path to the manpages and then export it. The manpages are in the `$IBRIXHOME/man` directory. For example, if `$IBRIXHOME` is `/usr/local/ibrix` (the default), you would set the `MANPATH` variable as follows on the management console and then export the variable.

```
MANPATH=$MANPATH:/usr/local/ibrix/man
```

## Changing passwords

You may want to change the passwords on your system:

- **Hardware passwords.** See the documentation for the specific hardware for more information.
- **Root password.** Use the `passwd(8)` command on each server in turn.
- **X9000 Software user password.** This password is created during installation and is used to log on to the management console GUI. The default is `ibrix`. You can change the password on the management console using the Linux `passwd` command. You will be prompted to enter the new password.

```
# passwd ibrix
```



---

## 3 Configuring virtual interfaces for client access

X9000 Software uses a cluster network interface to carry management console traffic and traffic between file serving nodes. This network is configured as `bond0` when the cluster is installed. For clusters with an agile management console configuration, a virtual interface is also created for the cluster network interface to provide failover support for the console.

Although the cluster network interface can carry traffic between file serving nodes and clients, HP recommends that you configure one or more user network interfaces for this purpose. Typically, `bond1` is created for the first user network when the cluster is configured.

To provide high availability for a user network, you should configure a bonded virtual interface (VIF) for the network and then set up failover for the VIF. This method prevents interruptions to client traffic. If necessary, the file serving node hosting the VIF can fail over to its standby backup node, and clients can continue to access the file system through the backup node.

### Network and VIF guidelines

To provide high availability, the user interfaces used for client access should be configured as bonded virtual interfaces (VIFs). Note the following:

- Nodes needing to communicate for file system coverage or for failover must be on the same network interface. Also, nodes set up as a failover pair must be connected to the same network interface.
- Use a Gigabit Ethernet port (or faster) for user networks.
- NFS, CIFS, FTP, and HTTP clients can use the same user VIF. The servers providing the VIF should be configured in backup pairs, and the NICs on those servers should also be configured for failover.
- For X9000 Linux and Windows clients, the servers hosting the VIF should be configured in backup pairs. However, X9000 clients do not support backup NICs. Instead, X9000 clients should connect to the parent bond of the user VIF or to a different VIF.

### Creating a bonded VIF

Use the following procedure to create a bonded VIF (`bond1:1` in this example):

1. If high availability (automated failover) is configured on the servers, disable it. Run the following command on the management console:

```
# ibrix_server -m -U
```

2. Identify the `bond1:1` VIF:

```
# ibrix_nic -a -n bond1:1 -h node1,node2,node3,node4
```

### 3. Assign an IP address to the bond1:1 VIFs on each node:

```
# ibrix_nic -c -n bond1:1 -h node1 -I 16.123.200.201 -M 255.255.255.0
# ibrix_nic -c -n bond1:1 -h node2 -I 16.123.200.202 -M 255.255.255.0
# ibrix_nic -c -n bond1:1 -h node3 -I 16.123.200.203 -M 255.255.255.0
# ibrix_nic -c -n bond1:1 -h node4 -I 16.123.200.204 -M 255.255.255.0
```

## Configuring standby backup nodes

Assign standby backup nodes for the bond1:1 interface. The backup nodes should be configured in pairs. For example, node1 is the backup for node2, and node2 is the backup for node1.

### 1. Identify the VIF:

```
# ibrix_nic -a -n bond1:2 -h node1,node2,node3,node4
```

### 2. Set up a standby server for each VIF:

```
# ibric_nic -b -H node1/bond1:1,node2/bond1:2
# ibric_nic -b -H node2/bond1:1,node1/bond1:2
# ibric_nic -b -H node3/bond1:1,node4/bond1:2
# ibric_nic -b -H node4/bond1:1,node3/bond1:2
```

## Configuring NIC failover

NIC monitoring should be configured on VIFs that will be used by NFS, CIFS, FTP, or HTTP. Use the same backup pairs that you used when configuring standby servers. For example:

```
# ibric_nic -m -h node1 -A node2/bond1:1
# ibric_nic -m -h node2 -A node1/bond1:1
# ibric_nic -m -h node3 -A node4/bond1:1
# ibric_nic -m -h node4 -A node3/bond1:1
```

## Configuring automated failover

To enable automated failover for your file serving nodes, execute the following command:

```
ibrix_server -m [-h SERVERNAME]
```

## Example configuration

This example uses two nodes, ib50-81 and ib50-82. These nodes are backups for each other, forming a backup pair.

```
[root@ib50-80 ~]# ibrix_server -l
Segment Servers
=====
```

SERVER_NAME	BACKUP	STATE	HA	ID	GROUP
ib50-81	ib50-82	Up	on	132cf61a-d25b-40f8-890e-e97363ae0d0b	servers
ib50-82	ib50-81	Up	on	7d258451-4455-484d-bf80-75c94d17121d	servers

All VIFs on ib50-81 have backup (standby) VIFs on ib50-82. Similarly, all VIFs on ib50-82 have backup (standby) VIFs on ib50-81. NFS, CIFS, FTP, and HTTP clients can connect to bond1:1 on either host. If necessary, the selected server will fail over to bond1:2 on the opposite host. X9000

clients could connect to `bond1` on either host, as these clients do not support or require NIC failover. (The following sample output shows only the relevant fields.)

```
[root@ib50-80 ~]# ibrix_nic -l
```

HOST	IFNAME	TYPE	STATE	IP_ADDRESS	BACKUP_HOST	BACKUP_IF
ib50-81	bond1:1	User	Up, LinkUp	16.226.50.220	ib50-82	bond1:1
ib50-81	bond0	Cluster	Up, LinkUp	172.16.0.81		
ib50-81	bond1:2	User	Inactive, Standby			
ib50-81	bond1	User	Up, LinkUp	16.226.50.81		
ib50-82	bond0	Cluster	Up, LinkUp	172.16.0.82		
ib50-82	bond1	User	Up, LinkUp	16.226.50.82		
ib50-82	bond1:2	User	Inactive, Standby			
ib50-82	bond1:1	User	Up, LinkUp	16.226.50.228	ib50-81	bond1:1

## Specifying VIFs in the client configuration

When you configure your clients, you may need to specify the VIF that should be used for client access.

**NFS/CIFS.** Specify the VIF IP address of the servers (for example, `bond1:0`) to establish connection. You can also configure DNS round robin to ensure NFS or CIFS client-to-server distribution. In both cases, the NFS/CIFS clients will cache the initial IP they used to connect to the respective share, usually until the next reboot.

**FTP.** When you add an FTP share on the Add FTP Shares dialog box or with the `ibrix_ftpshare` command, specify the VIF as the IP address that clients should use to access the share.

**HTTP.** When you create a virtual host on the Create Vhost dialog box or with the `ibrix_httpvhost` command, specify the VIF as the IP address that clients should use to access shares associated with the Vhost.

**X9000 clients.** Use the following command to prefer the appropriate user network. Execute the command once for each destination host that the client should contact using the specified interface.

```
ibrix_client -n -h SRCHOST -A DESTNOST/IFNAME
```

For example:

```
ibrix_client -n -h client12.mycompany.com -A ib50-81.mycompany.com/bond1
```



### NOTE:

Because the backup NIC cannot be used as a preferred network interface for X9000 clients, add one or more user network interfaces to ensure that HA and client communication work together.





---

## 4 Configuring failover

This chapter describes how to configure failover for agile management consoles, file serving nodes, network interfaces, and HBAs.

### Agile management consoles

The management console maintains the cluster configuration and provides graphical and command-line user interfaces for managing and monitoring the cluster. Typically, one active management console and one passive management console are installed when the cluster is installed. This is called an *agile* management console configuration.



---

#### NOTE:

Optionally, the management console can be installed on a dedicated Management Server. This section describes the agile management console configuration.

---

### Agile management console modes

An agile management console can be in one of the following modes:

- **active.** In this mode, the management console controls console operations. All cluster administration and configuration commands must be run from the active management console.
- **passive.** In this mode, the management console monitors the health of the active management console. If the active management console fails, the passive management console becomes the active console.
- **maintenance.** In this mode, the management console does not participate in console operations. Maintenance mode should be used for operations such as manual failover of the active management console, X9000 Software upgrades, and blade replacements.

### Agile management consoles and failover

Using an agile management console configuration provides high availability for management console services. If the active management console fails, the cluster virtual interface will go down. When the passive management console detects that the cluster virtual interface is down, it will become the active console. This management console rebuilds the cluster virtual interface, starts management console services locally, transitions into active mode, and take over management console operation.

Failover of the active management console affects the following features:

- **User networks.** The virtual interface used by clients will also fail over. Users may notice a brief reconnect while the newly active management console takes over management of the virtual interface.
- **Support tickets.** The existing support ticket information is not moved to the newly active management console. Support Ticket operations are always handled by the active management console and the final output of the operations is stored there.

- **Management console GUI.** You will need to reconnect to the management console VIF after the failover.

## Failing over the management console manually

To fail over the active management console manually, place the console into maintenance mode. Enter the following command on the node hosting the console:

```
ibrix_fm -m maintenance
```

The command takes effect immediately.

The failed-over management console remains in maintenance mode until it is moved to passive mode using the following command:

```
ibrix_fm -m passive
```

A management console cannot be moved from maintenance mode to active mode.

## Viewing information about management consoles

To view mode information, use the following command:

```
ibrix_fm -i
```



### NOTE:

If the management console was not installed in an agile configuration, the output will report FusionServer: fusion manager name not set! (active, quorum is not configured).

When a management console is installed, it is registered in the management console configuration. To view a list of all registered management consoles, use the following command:

```
ibrix_fm -f
```

## Cluster high availability

X9000 Software High Availability keeps your data accessible at all times. Failover protection can be configured for file serving nodes, network interfaces, individual segments, and HBAs. Through physical and logical configuration policies, you can set up a flexible and scalable high availability solution. X9000 clients experience no changes in service and are unaware of the failover events.

## Failover modes

High Availability has two failover modes: the default *manual failover* and the optional *automated failover*. A manual failover uses the `ibrix_server` command or the management console GUI to fail over a file serving node to its standby. The server can be powered down or remain up during the procedure. Manual failover also includes failover of any network interfaces having defined standbys. You can perform a manual failover at any time, regardless of whether automated failover is in effect.

Automated failover allows the management console to initiate failover when it detects that standby-protected components have failed. A basic automated failover setup protects all file serving nodes. A comprehensive setup also includes network interface monitoring to protect user network interfaces and HBA monitoring to protect access from file serving nodes to storage via an HBA.

When automated failover is enabled, the management console listens for heartbeat messages that the file serving nodes broadcast at one-minute intervals. The management console automatically initiates failover when it fails to receive five consecutive heartbeats or, if HBA monitoring is enabled, when a heartbeat message indicates that a monitored HBA or pair of HBAs has failed.

If network interface monitoring is enabled, automated failover occurs when the management console receives a heartbeat message indicating that a monitored network might be down and then the console cannot reach that interface.

If a file serving node fails over, you will need to manually fail back the node.

## What happens during a failover

The following events occur during automated or manual failover of a file serving node to its standby:

1. The management console verifies that the standby is powered on and accessible.
2. The management console migrates ownership of the node's segments to the standby and notifies all file serving nodes and X9000 clients about the migration. This is a persistent change.
3. If network interface monitoring has been set up, the management console activates the standby user network interface and transfers the IP address of the node's user network interface to it.

To determine the progress of a failover, view the Status tab on the GUI or execute the `ibrix_server -l` command. While the management console is migrating segment ownership, the operational status of the node is Up-InFailover or Down-InFailover, depending on whether the node was powered up or down when failover was initiated. When failover is complete, the operational status changes to Up-FailedOver or Down-FailedOver. For more information about operational states, see ["Monitoring the status of file serving nodes"](#) on page 51.

Both automated and manual failovers trigger an event that is reported on the GUI.

## Setting up automated failover

The recommended minimum setup for automated failover protection is as follows:

1. Identify standbys for file serving nodes or specific segments. You must implement either server-level or segment-level standby protection; you cannot implement both.
2. Identify power sources for file serving nodes. For APC power sources, associate file serving nodes to power source slots.
3. Turn on automated failover.

If your cluster includes one or more user network interfaces carrying NFS/CIFS client traffic, HP recommends that you identify standby network interfaces and set up network interface monitoring.

If your file serving nodes are connected to storage via HBAs, HP recommends that you set up HBA monitoring.

## Identifying standbys for file serving nodes

file serving nodes can be configured to provide standby service for one another in the following configurations:

- **1 x 1.** Set up standby pairs, where each server in a pair is the standby for the other.
- **1 x N.** Assign the same standby to a certain number of primaries.

Contact HP Support for recommendations based on your environment.

The following restrictions apply to all types of standby configurations:

- The management console must have access to both the primary server and its standby.
- The same file system must be mounted on both the primary server and its standby.
- A server identified as a standby must be able to see all segments that might fail over to it.
- In a SAN environment, a primary server and its standby must use the same storage infrastructure to access a segment's physical volumes (for example, a multiported RAID array).

To identify a standby for a file serving node, use the following command:

```
<installdirectory>/bin/ibrix_server -b -h HOSTNAME1,HOSTNAME2
```

For example, to identify node s2.hp.com as the standby for all segments on node s1.hp.com:

```
<installdirectory>/bin/ibrix_server -b -h s1.hp.com,s2.hp.com
```

For performance reasons, you might want to fail over specific segments to a standby instead of failing over all segments on a node to a standby. Use this command to identify the segments:

```
<installdirectory>/bin/ibrix_fs -b -f FSNAME -s LVLIST -h HOSTNAME
```

For example, to identify node s1.hp.com as the standby for segments ilv\_1, ilv\_2, and ilv\_3 in file system ifs1:

```
<installdirectory>/bin/ibrix_fs -b -f ifs1 -s ilv_1,ilv_2,ilv_3 -h s1.hp.com
```

## Identifying power sources

To implement automated failover, perform a forced manual failover, or remotely power a file serving node up or down, you must set up programmable power sources for the nodes and their standbys. Using programmable power sources prevents a “split-brain scenario” between a failing file serving node and its standby, allowing the failing server to be centrally powered down by the management console in the case of automated failover, and manually in the case of a forced manual failover.

X9000 Software works with iLO, IPMI, OpenIPMI, and OpenIPMI2 integrated power sources and with APC power sources.

### Preliminary configuration

Certain configuration steps are required when setting up power sources:

- **All types.** If you plan to implement automated failover, ensure that the management console has LAN access to the power sources.
- **Integrated power sources.** Install the environment and any drivers and utilities, as specified by the vendor documentation. If you plan to protect access to the power sources, set up the UID and password to be used.
- **APC.** Enable SNMP access. Set the Community Name to `ibrix` and the Access Type to `write+`. If `write+` does not work with your configuration, set the Access Type to `write`.

### Identifying power sources

All power sources must be identified to the configuration database before they can be used.

**Integrated power sources.** To identify an integrated power source, use the following command:

```
<installdirectory>/bin/ibrix_powersrc -a -t {ipmi|openipmi|openipmi2|iLO}
-h HOSTNAME -I IPADDR -u USERNAME -p PASSWORD
```

For example, to identify an iLO power source at IP address 192.168.3.170 for node ss01:

```
<installdirectory>/bin/ibrix_powersrc -a -t ilo -h ss01 -I 192.168.3.170
-u Administrator -p password
```

**APC power source.** To identify an APC power source, use the following command:

```
<installdirectory>/bin/ibrix_powersrc -a -t {apc|apc_msp} -h POWERSRCNAME -n NUMSLOTS  
-I IPADDR
```

For example, to identify an eight-port APC power source named `ps1` at IP address 192.168.3.150:

```
<installdirectory>/bin/ibrix_powersrc -a -t apc -h ps1 -n 8 -I 192.168.3.150
```

For APC power sources, you must also associate file serving nodes to power source slots. (This step is unnecessary for integrated power sources because the nodes are connected by default to slot 1.) Use the following command:

```
<installdirectory>/bin/ibrix_hostpower -a -i SLOTID -s POWERSOURCE -h HOSTNAME
```

For example, to identify that node `s1.hp.com` is connected to slot 1 on APC power source `ps1`:

```
<installdirectory>/bin/ibrix_hostpower -a -i 1 -s ps1 -h s1.hp.com
```

### Updating the configuration database with power source changes

If you move a file serving node to a different power source slot, unplug it from a power source slot, or change its IP address or password, you must update the configuration database with the changes. To do this, use the following command. The user name and password options are needed only for remotely managed power sources. Include the `-s` option to have the management console skip BMC.

```
<installdirectory>/bin/ibrix_powersrc -m [-I IPADDR] [-u USERNAME] [-p PASSWORD]  
[-s] -h POWERSRCLIST
```

The following command changes the IP address for power source `ps1`:

```
<installdirectory>/bin/ibrix_powersrc -m -I 192.168.3.153 -h ps1
```

To change the APC slot association for a file serving node, use the following command:

```
<installdirectory>/bin/ibrix_hostpower -m -i FROM_SLOT_ID,TO_SLOT_ID -s POWERSOURCE  
-h HOSTNAME
```

For example, to identify that node `s1.hp.com` has been moved from slot 3 to slot 4 on APC power source `ps1`:

```
<installdirectory>/bin/ibrix_hostpower -m -i 3,4 -s ps1 -h s1.hp.com
```

### Dissociating a file serving node from a power source

You can dissociate a file serving node from an integrated power source by dissociating it from slot 1 (its default association) on the power source. Use the following command:

```
<installdirectory>/bin/ibrix_hostpower -d -s POWERSOURCE -h HOSTNAME
```

To dissociate a file serving node from an APC power source on the specified slot, use the following command. To dissociate the node from all slots on the power source, omit the `-i` option.

```
<installdirectory>/bin/ibrix_hostpower -d [-s POWERSOURCE [-i SLOT]] -h HOSTNAME
```

For example, to dissociate file serving node `s1.hp.com` from slot 3 on APC power source `ps1`:

```
<installdirectory>/bin/ibrix_hostpower -d -s ps1 -i 3 -h s1.hp.com
```

### Deleting power sources from the configuration database

To conserve storage, delete power sources that are no longer in use from the configuration database. If you are deleting multiple power sources, use commas to separate them.

```
<installdirectory>/bin/ibrix_powersrc -d -h POWERSRCLIST
```

## Turning automated failover on and off

Automated failover is turned off by default. When automated failover is turned on, the management console starts monitoring heartbeat messages from file serving nodes. You can turn automated failover on and off for all file serving nodes or for selected nodes.

To turn on automated failover, use the following command:

```
<installdirectory>/bin/ibrix_server -m [-h SERVERNAME]
```

To turn off automated failover, include the `-U` option:

```
<installdirectory>/bin/ibrix_server -m -U [-h SERVERNAME]
```

To turn automated failover on or off for a single file serving node, include the `-h SERVERNAME` option.

## Manually failing over a file serving node

To set up a cluster for manual failover, first identify server-level or segment-level standbys for each file serving node, as described in “[Identifying standbys for file serving nodes](#)” on page 27.

Manual failover does not require the use of programmable power supplies. However, if you have installed and identified power supplies for file serving nodes, you can power down a server before manually failing it over. You can fail over a file serving node manually, even when automated failover is turned on.

A file serving node can be failed over from the GUI or the CLI.

On the CLI, complete the following steps:

1. Run `ibrix_server -f`, specifying the node to be failed over in the `HOSTNAME` option. If appropriate, include the `-p` option to power down the node before segments are migrated:

```
<installdirectory>/bin/ibrix_server -f [-p] -h HOSTNAME
```

2. Determine whether the failover was successful:

```
<installdirectory>/bin/ibrix_server -l
```

The contents of the `STATE` field indicate the status of the failover. If the field persistently shows `Down-InFailover` or `Up-InFailover`, the failover did not complete; contact HP Support for assistance. For information about the values that can appear in the `STATE` field, see “[What happens during a failover](#)” on page 27.

## Failing back a file serving node

After automated or manual failover of a file serving node, you must manually fail back the server, which restores ownership of the failed-over segments and network interfaces to the server. Before failing back the node, confirm that the primary server can see all of its storage resources and networks. The segments owned by the primary server will not be accessible if the server cannot see its storage.

To fail back a file serving node, use the following command. The `HOSTNAME` argument specifies the name of the failed-over node.

```
<installdirectory>/bin/ibrix_server -f -U -h HOSTNAME
```

After failing back the node, determine whether the failback completed fully. If the failback is not complete, contact HP Support for assistance.



#### NOTE:

A failback might not succeed if the time period between the failover and the failback is too short, and the primary server has not fully recovered. HP recommends ensuring that both servers are up and running and then waiting 60 seconds before starting the failback. Use the `ibrix_server -l` command to verify that the primary server is up and running. The status should be Up-FailedOver before performing the failback.

## Using network interface monitoring

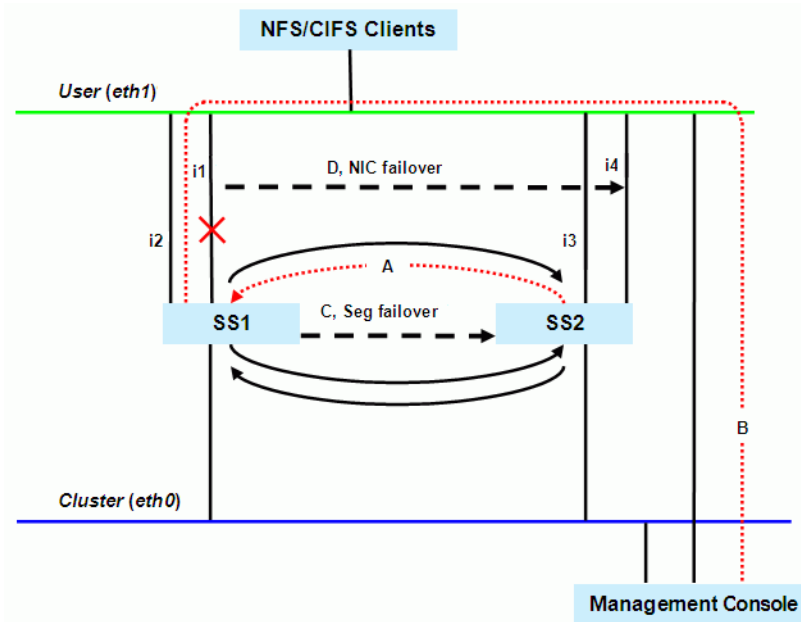
With network interface monitoring, one file serving node monitors another file serving node over a designated network interface. If the monitoring server loses contact with its destination server over the interface, it notifies the management console. If the management console also cannot contact the destination server over that interface, it fails over both the destination server and the network interface to their standbys. Clients that were mounted on the failed-over server do not experience any service interruption and are unaware that they are now mounting the file system on a different server.

Unlike X9000 clients, NFS and CIFS clients cannot reroute file requests to a standby if the file serving node where they are mounted should fail. To ensure continuous client access to files, HP recommends that you put NFS/CIFS traffic on a user network interface (see [“Preferring network interfaces”](#) on page 65), and then implement network interface monitoring for it.

Comprehensive protection of NFS/CIFS traffic also involves setting up network interface monitoring for the cluster interface. Although the management console will eventually detect interruption of a file serving node’s connection to the cluster interface and initiate segment failover if automated failover is turned on, failover will occur much faster if the interruption is detected via network interface monitoring. (If automated failover is not turned on, you will begin to see file access problems if the cluster interface fails.) There is no difference in the way that monitoring is set up for the cluster interface and a user network interface. In both cases, you set up file serving nodes to monitor each other over the interface.

## Sample scenario

The following diagram illustrates a monitoring and failover scenario in which a 1:1 standby relationship is configured. Each standby pair is also a network interface monitoring pair. When SS1 loses its connection to the user network interface (`eth1`), as shown by the red X, SS2 can no longer contact SS1 (A). SS2 notifies the management console, which then tests its own connection with SS1 over `eth1` (B). The management console cannot contact SS1 on `eth1`, and initiates failover of SS1’s segments (C) and user network interface (D).



## Identifying standbys

To protect a network interface, you must identify a standby for it on each file serving node that connects to the interface. The following restrictions apply when identifying a standby network interface:

- The standby network interface must be unconfigured and connected to the same switch (network) as the primary interface.
- The file serving node that supports the standby network interface must have access to the file system that the clients on that interface will mount.

Virtual interfaces are highly recommended for handling user network interface failovers. If a VIF user network interface is teamed/bonded, failover occurs only if all teamed network interfaces fail. Otherwise, traffic switches to the surviving teamed network interfaces.

To identify standbys for a network interface, execute the following command once for each file serving node. *IFNAME1* is the network interface that you want to protect and *IFNAME2* is the standby interface.

```
<installdirectory>/bin/ibrix_nic -b -H HOSTNAME1/IFNAME1,HOSTNAME2/IFNAME2
```

The following command identifies virtual interface `eth2:2` on file serving node `s2.hp.com` as the standby interface for interface `eth2` on file serving node `s1.hp.com`:

```
<installdirectory>/bin/ibrix_nic -b -H s1.hp.com/eth2,s2.hp.com/eth2:2
```

## Setting up a monitor

File serving node failover pairs can be identified as network interface monitors for each other. Because the monitoring must be declared in both directions, this is a two-pass process for each failover pair.

To set up a network interface monitor, use the following command:

```
<installdirectory>/bin/ibrix_nic -m -h MONHOST -A DESTHOST/IFNAME
```

For example, to set up file serving node `s2.hp.com` to monitor file serving node `s1.hp.com` over user network interface `eth1`:

```
<installdirectory>/bin/ibrix_nic -m -h s2.hp.com -A s1.hp.com/eth1
```

To delete network interface monitoring, use the following command:



```
<installdirectory>/bin/ibrix_nic -m -h MONHOST -D DESTHOST/IFNAME
```

## Deleting standbys

To delete a standby for a network interface, use the following command:

```
<installdirectory>/bin/ibrix_nic -b -U HOSTNAME1/IFNAME1
```

For example, to delete the standby that was assigned to interface `eth2` on file serving node `s1.hp.com`:

```
<installdirectory>/bin/ibrix_nic -b -U s1.hp.com/eth2
```

## Setting up HBA monitoring

You can configure High Availability to initiate automated failover upon detection of a failed HBA. HBA monitoring can be set up for either dual-port HBAs with built-in standby switching or single-port HBAs, whether standalone or paired for standby switching via software. X9000 Software does not play any role in vendor- or software-mediated HBA failover—traffic moves to the remaining functional port without any management console involvement.

HBAs use worldwide names for some parameter values. These are either worldwide node names (WWNN) or worldwide port names (WWPN). The WWPN is the name an HBA presents when logging in to a SAN fabric. Worldwide names consist of 16 hexadecimal digits grouped in pairs. In X9000 Software, these are written as dot-separated pairs (for example, 21.00.00.e0.8b.05.05.04).

To set up HBA monitoring, first discover the HBAs, and then perform the procedure that matches your HBA hardware:

- For single-port HBAs without built-in standby switching: Turn on HBA monitoring for all ports that you want to monitor for failure (see [“Turning HBA monitoring on or off”](#) on page 34).
- For dual-port HBAs with built-in standby switching and single-port HBAs that have been set up as standby pairs via software: Identify the standby pairs of ports to the configuration database (see [“Identifying standby-paired HBA ports”](#) on page 34, and then turn on HBA monitoring for all paired ports (see [“Turning HBA monitoring on or off”](#) on page 34). If monitoring is turned on for just one port in a standby pair and that port then fails, the management console will fail over the server even though the HBA has automatically switched traffic to the surviving port. When monitoring is turned on for both ports, the management console initiates failover only when both ports in a pair fail.

When both HBA monitoring and automated failover for file serving nodes are set up, the management console will fail over a server in two situations:

- Both ports in a monitored set of standby-paired ports fail. Because, during the HBA monitoring setup, all standby pairs were identified in the configuration database, the management console knows that failover is required only when both ports fail.
- A monitored single-port HBA fails. Because no standby has been identified for the failed port, the management console knows to initiate failover immediately.

## Discovering HBAs

You must discover HBAs before you set up HBA monitoring, when you replace an HBA, and when you add a new HBA to the cluster. Discovery informs the configuration database of only a port's WWPN. You must identify ports that are teamed as standby pairs. Use the following command:

```
<installdirectory>/bin/ibrix_hba -a [-h HOSTLIST]
```

## Identifying standby-paired HBA ports

Identifying standby-paired HBA ports to the configuration database allows the management console to apply the following logic when they fail:

- If one port in a pair fails, do nothing. Traffic will automatically switch to the surviving port, as configured by the vendor or the software.
- If both ports in a pair fail, fail over the server's segments to the standby server.

Use the following command to identify two HBA ports as a standby pair:

```
<installdirectory>/bin/ibrix_hba -b -P WWP1:WWP2 -h HOSTNAME
```

Enter the WWPN as decimal-delimited pairs of hexadecimal digits. The following command identifies port 20.00.12.34.56.78.9a.bc as the standby for port 42.00.12.34.56.78.9a.bc for the HBA on file serving node s1.hp.com:

```
<installdirectory>/bin/ibrix_hba -b -P 20.00.12.34.56.78.9a.bc:42.00.12.34.56.78.9a.bc  
-h s1.hp.com
```

## Turning HBA monitoring on or off

If your cluster uses single-port HBAs, turn on monitoring for all of the ports to set up automated failover in the event of HBA failure. Use the following command:

```
<installdirectory>/bin/ibrix_hba -m -h HOSTNAME -p PORT
```

For example, to turn on HBA monitoring for port 20.00.12.34.56.78.9a.bc on node s1.hp.com:

```
<installdirectory>/bin/ibrix_hba -m -h s1.hp.com -p 20.00.12.34.56.78.9a.bc
```

To turn off HBA monitoring for an HBA port, include the `-U` option:

```
<installdirectory>/bin/ibrix_hba -m -U -h HOSTNAME -p PORT
```

## Deleting standby port pairings

Deleting port pairing information from the configuration database does not remove the standby pairing of the ports. The standby pairing is either built in by the vendor or implemented by software.

To delete standby-paired HBA ports from the configuration database, enter the following command:

```
<installdirectory>/bin/ibrix_hba -b -U -P WWP1:WWP2 -h HOSTNAME
```

For example, to delete the pairing of ports 20.00.12.34.56.78.9a.bc and 42.00.12.34.56.78.9a.bc on node s1.hp.com:

```
<installdirectory>/bin/ibrix_hba -b -U -P 20.00.12.34.56.78.9a.bc:42.00.12.34.56.78.9a.bc  
-h s1.hp.com
```

## Deleting HBAs from the configuration database

Before switching an HBA card to a different machine, delete the HBA from the configuration database. Use the following command:

```
<installdirectory>/bin/ibrix_hba -d -h HOSTNAME -w WWNN
```

## Displaying HBA information

Use the following command to view information about the HBAs in the cluster. To view information for all hosts, omit the `-h HOSTLIST` argument.

```
<installdirectory>/bin/ibrix_hba -l [-h HOSTLIST]
```

The following table describes the fields in the output.

Field	Description
Host	Server on which the HBA is installed.
Node WWN	This HBA's WWNN.
Port WWN	This HBA's WWPN.
Port State	Operational state of the port.
Backup Port WWN	WWPN of the standby port for this port (standby-paired HBAs only).
Monitoring	Whether HBA monitoring is enabled for this port.

## Checking the High Availability configuration

Use the `ibrix_haconfig` command to determine whether High Availability features have been configured for specific file serving nodes. The command checks for the following features and provides either a summary or a detailed report of the results:

- Programmable power source
- Standby server or standby segments
- Cluster and user network interface monitors
- Standby network interface for each user network interface
- HBA port monitoring
- Whether automated failover is on

For each High Availability feature, the summary report returns one of the following results for each tested file serving node and optionally for their standbys:

- **Passed.** The feature has been configured.
- **Warning.** The feature has not been configured, but the significance of the finding is not clear. For example, the absence of discovered HBAs can indicate either that the HBA monitoring feature was not configured or that HBAs are not physically present on the tested servers.
- **Failed.** The feature has not been configured.

The detailed report includes an overall result status for all tested file serving nodes and describes details about the checks performed on each High Availability feature. By default, the report includes details only about checks that received a Failed or a Warning result. You can expand the report to include details about checks that received a Passed result.

## Viewing a summary report

Executing the `ibrix_haconfig` command with no arguments returns a summary of all file serving nodes. To check specific file serving nodes, include the `-h HOSTLIST` argument. To check standbys, include the `-b` argument. To view results only for file serving nodes that failed a check, include the `-f` argument.

```
<installdirectory>/bin/ibrix_haconfig -l [-h HOSTLIST] [-f] [-b]
```

For example, to view a summary report for file serving nodes `xs01.hp.com` and `xs02.hp.com`:

```
<installdirectory>/bin/ibrix_haconfig -l -h xs01.hp.com,xs02.hp.com
```

Host	HA Configuration	Power Sources	Backup Servers	Auto Failover
Nics Monitored	Standby Nics	HBAs Monitored		
xs01.hp.com	FAILED	PASSED	PASSED	PASSED
FAILED	PASSED	FAILED		
xs02.hp.com	FAILED	PASSED	FAILED	FAILED
FAILED	WARNED	WARNED		

## Viewing a detailed report

Execute the `ibrix_haconfig -i` command to view the detailed report:

```
<installdirectory>/bin/ibrix_haconfig -i [-h HOSTLIST] [-f] [-b] [-s] [-v]
```

The `-h HOSTLIST` option lists the nodes to check. To also check standbys, include the `-b` option. To view results only for file serving nodes that failed a check, include the `-f` argument. The `-s` option expands the report to include information about the file system and its segments. The `-v` option produces detailed information about configuration checks that received a Passed result.

For example, to view a detailed report for file serving nodes `xs01.hp.com`:

```
<installdirectory>/bin/ibrix_haconfig -i -h xs01.hp.com
```

```
----- Overall HA Configuration Checker Results -----
FAILED
----- Overall Host Results -----
Host          HA Configuration Power Sources Backup Servers Auto Failover Nics Monitored
      Standby Nics  HBAs Monitored
xs01.hp.com FAILED          PASSED          PASSED          PASSED          FAILED
      PASSED          FAILED
----- Server xs01.hp.com FAILED Report -----

Check Description                                     Result  Result Information
=====
Power source(s) configured                             PASSED
Backup server or backups for segments configured PASSED
Automatic server failover configured                   PASSED

Cluster & User Nics monitored
  Cluster nic xs01.hp.com/eth1 monitored               FAILED  Not monitored

User nics configured with a standby nic                PASSED

HBA ports monitored
  Hba port 21.01.00.e0.8b.2a.0d.6d monitored           FAILED  Not monitored
  Hba port 21.00.00.e0.8b.0a.0d.6d monitored           FAILED  Not monitored
```

---

# 5 Configuring cluster event notification

## Setting up email notification of cluster events

You can set up event notifications by event type or for one or more specific events. To set up automatic email notification of cluster events, associate the events with email recipients and then configure email settings to initiate the notification process.

### Associating events and email addresses

You can associate any combination of cluster events with email addresses: all Alert, Warning, or Info events, all events of one type plus a subset of another type, or a subset of all types.

The notification threshold for Alert events is 90% of capacity. Threshold-triggered notifications are sent when a monitored system resource exceeds the threshold and are reset when the resource utilization dips 10% below the threshold. For example, a notification is sent the first time usage reaches 90% or more. The next notice is sent only if the usage declines to 80% or less (event is reset), and subsequently rises again to 90% or above.

To associate all types of events with recipients, omit the `-e` argument in the following command. Use the `ALERT`, `WARN`, and `INFO` keywords to make specific type associations or use `EVENTLIST` to associate specific events.

```
<installdirectory>/bin/ibrix_event -c [-e ALERT|WARN|INFO|EVENTLIST] -m EMAILLIST
```

The following command associates all types of events to `admin@hp.com`:

```
<installdirectory>/bin/ibrix_event -c -m admin@hp.com
```

The next command associates all Alert events and two Info events to `admin@hp.com`:

```
<installdirectory>/bin/ibrix_event -c -e ALERT,server.registered,filesystem.space.full  
-m admin@hp.com
```

### Configuring email notification settings

Configuring email notification settings involves specifying the SMTP server and header information and turning the notification process on or off. The state of the email notification process has no effect on the display of cluster events in the management console GUI.

The server must be able to receive and send email and must recognize the From and Reply-to addresses. Be sure to specify valid email addresses, especially for the SMTP server. If an address is not valid, the SMTP server will reject the email.

```
<installdirectory>/bin/ibrix_event -m on|off -s SMTP -f from [-r reply-to] [-t subject]
```

The following command configures email settings to use the `mail.hp.com` SMTP server and to turn on notifications:

```
<installdirectory>/bin/ibrix_event -m on -s mail.hp.com -f FM@hp.com  
-r MIS@hp.com -t Cluster1 Notification
```

## Turning email notifications on or off

After configuration is complete, use the `-m on` option to turn on email notifications. To turn off email notifications, use the `-m off` option.

```
<installdirectory>/bin/ibrix_event -m on|off -s SMTP -f from
```

## Dissociating events and email addresses

To remove the association between events and email addresses, use the following command:

```
<installdirectory>/bin/ibrix_event -d [-e ALERT|WARN|INFO|EVENTLIST] -m EMAILLIST
```

For example, to dissociate event notifications for `admin@hp.com`:

```
<installdirectory>/bin/ibrix_event -d -m admin@hp.com
```

To turn off all Alert notifications for `admin@hp.com`:

```
<installdirectory>/bin/ibrix_event -d -e ALERT -m admin@hp.com
```

To turn off the `server.registered` and `filesystem.created` notifications for `admin1@hp.com` and `admin2@hp.com`:

```
<installdirectory>/bin/ibrix_event -d -e server.registered,filesystem.created  
-m admin1@hp.com,admin2@hp.com
```

## Testing email addresses

To test an email address with a test message, notifications must be turned on. If the address is valid, the command signals success and sends an email containing the settings to the recipient. If the address is not valid, the command returns an `address failed` exception.

```
<installdirectory>/bin/ibrix_event -u -n EMAILADDRESS
```

## Viewing email notification settings

The `ibrix_event` command provides comprehensive information about email settings and configured notifications.

```
<installdirectory>/bin/ibrix_event -L
```

Sample output follows:

```
Email Notification : Enabled  
SMTP Server       : mail.hp.com  
From              : FM@hp.com  
Reply To          : MIS@hp.com
```

EVENT	LEVEL	TYPE	DESTINATION
-----	-----	-----	-----
asyncrep.completed	ALERT	EMAIL	admin@hp.com
asyncrep.failed	ALERT	EMAIL	admin@hp.com

## Setting up SNMP notifications

X9000 Software supports SNMP (Simple Network Management Protocol) V1, V2, and V3.



#### NOTE:

Users of software versions earlier than 4.3 should be aware that the single `ibrix_snmp` command has been replaced by two commands, `ibrix_snmpagent` and `ibrix_snmptrap`. If you have scripts that include `ibrix_snmp`, be sure to edit them to include the correct commands.

Whereas SNMPV2 security was enforced by use of community password strings, V3 introduces the USM and VACM. Discussion of these models is beyond the scope of this document. Refer to RFCs 3414 and 3415 at <http://www.ietf.org> for more information. Note the following:

- In the SNMPV3 environment, every message contains a user name. The function of the USM is to authenticate users and ensure message privacy through message encryption and decryption. Both authentication and privacy, and their passwords, are optional and will use default settings where security is less of a concern.
- With users validated, the VACM determines which managed objects these users are allowed to access. The VACM includes an access scheme to control user access to managed objects; context matching to define which objects can be accessed; and MIB views, defined by subsets of IOD subtree and associated bitmask entries, which define what a particular user can access in the MIB.

Steps for setting up SNMP include:

- Agent configuration (all SNMP versions)
- Trapsink configuration (all SNMP versions)
- Associating event notifications with trapsinks (all SNMP versions)
- View definition (V3 only)
- Group and user configuration (V3 only)

X9000 Software implements an SNMP agent on the management console that supports the private X9000 Software MIB. The agent can be polled and can send SNMP traps to configured trapsinks.

Setting up SNMP notifications is similar to setting up email notifications. You must associate events to trapsinks and configure SNMP settings for each trapsink to enable the agent to send a trap when an event occurs.

## Configuring the SNMP agent

The SNMP agent is created automatically when the management console is installed. It is initially configured as an SNMPv2 agent and is off by default.

Some SNMP parameters and the SNMP default port are the same, regardless of SNMP version. The agent port is 161 by default. *SYSCONTACT*, *SYSNAME*, and *SYSLOCATION* are optional MIB-II agent parameters that have no default values.

The `-c` and `-s` options are also common to all SNMP versions. The `-c` option turns the encryption of community names and passwords on or off. There is no encryption by default. Using the `-s` option toggles the agent on and off; it turns the agent on by starting a listener on the SNMP port, and turns it off by shutting off the listener. The default is off.

The format for a v1 or v2 update command follows:

```
ibrix_snmpagent -u -v {1|2} [-p PORT] [-r READCOMMUNITY] [-w WRITECOMMUNITY]
[-t SYSCONTACT] [-n SYSNAME] [-o SYSLOCATION] [-c {yes|no}] [-s {on|off}]
```

The update command for SNMPv1 and v2 uses optional community names. By convention, the default `READCOMMUNITY` name used for read-only access and assigned to the agent is `public`. No default `WRITECOMMUNITY` name is set for read-write access (although the name `private` is often used).

The following command updates a v2 agent with the write community name `private`, the agent's system name, and that system's physical location:

```
ibrix_snmpagent -u -v 2 -w private -n agenthost.domain.com -o DevLab-B3-U6
```

The SNMPv3 format adds an optional `engine id` that overrides the default value of the agent's host name. The format also provides the `-y` and `-z` options, which determine whether a v3 agent can process v1/v2 read and write requests from the management station. The format is:

```
ibrix_snmpagent -u -v 3 [-e engineId] [-p PORT] [-r READCOMMUNITY]
[-w WRITECOMMUNITY] [-t SYSCONTACT] [-n SYSNAME] [-o SYSLOCATION]
[-y {yes|no}] [-z {yes|no}] [-c {yes|no}] [-s {on|off}]
```

## Configuring trapsink settings

A *trapsink* is the host destination where agents send *traps*, which are asynchronous notifications sent by the agent to the management station. A trapsink is specified either by name or IP address. X9000 Software supports multiple trapsinks; you can define any number of trapsinks of any SNMP version, but you can define only one trapsink per host, regardless of the version.

At a minimum, trapsink configuration requires a destination host and SNMP version. All other parameters are optional and many assume the default value if no value is specified. Trapsink configuration for SNMPv3 is more detailed than for earlier versions. The main differences involve the additional security parameters added by SNMPv3.

The format for creating a v1/v2 trapsink is:

```
ibrix_snmptrap -c -h HOSTNAME -v {1|2} [-p PORT] [-m COMMUNITY] [-s {on|off}]
```

If a port is not specified, the command defaults to port 162. If a community is not specified, the command defaults to the community name `public`. The `-s` option toggles agent trap transmission on and off. The default is on. For example, to create a v2 trapsink with a new community name, enter:

```
ibrix_snmptrap -c -h lab13-116 -v 2 -m private
```

For a v3 trapsink, additional options define security settings. `USERNAME` is a v3 user defined on the trapsink host and is required. The security level associated with the trap message depends on which passwords are specified—the authentication password, both the authentication and privacy passwords, or no passwords. The `CONTEXT_NAME` is required if the trap receiver has defined subsets of managed objects. The format is:

```
ibrix_snmptrap -c -h HOSTNAME -v 3 [-p PORT] -n USERNAME [-j {MD5|SHA}]
[-k AUTHORIZATION_PASSWORD] [-y {DES|AES}] [-z PRIVACY_PASSWORD]
[-x CONTEXT_NAME] [-s {on|off}]
```

The following command creates a v3 trapsink with a named user and specifies the passwords to be applied to the default algorithms. If specified, passwords must contain at least eight characters.

```
ibrix_snmptrap -c -h lab13-114 -v 3 -n trapsender -k auth-passwd -z priv-passwd
```

## Associating events and trapsinks

Associating events with trapsinks is similar to associating events with email recipients, except that you specify the host name or IP address of the trapsink instead of an email address.

Use the `ibrix_event` command to associate SNMP events with trapsinks. The format is:



```
<installdirectory>/bin/ibrix_event -c -y SNMP [-e ALERT|INFO|EVENTLIST]
-m TRAPSINK
```

For example, to associate all Alert events and two Info events with a trapsink at IP address 192.168.2.32, enter:

```
<installdirectory>/bin/ibrix_event -c -y SNMP -e ALERT,server.registered,
filesystem.created -m 192.168.2.32
```

Use the `ibrix_event -d` command to dissociate events and trapsinks:

```
<installdirectory>/bin/ibrix_event -d -y SNMP [-e ALERT|INFO|EVENTLIST] -m TRAPSINK
```

## Defining views

A MIB view is a collection of paired OID subtrees and associated bitmasks that identify which subidentifiers are significant to the view's definition. Using the bitmasks, individual OID subtrees can be included in or excluded from the view.

An instance of a managed object belongs to a view if:

- The OID of the instance has at least as many sub-identifiers as the OID subtree in the view.
- Each sub-identifier in the instance and the subtree match when the bitmask of the corresponding sub-identifier is nonzero.

The management console automatically creates the `excludeAll` view that blocks access to all OIDs. This view cannot be deleted; it is the default read and write view if one is not specified for a group with the `ibrix_snmpgroup` command. The catch-all OID and mask are:

```
OID = .1
Mask = .1
```

Consider these examples, where instance `.1.3.6.1.2.1.1` matches, instance `.1.3.6.1.4.1` matches, and instance `.1.2.6.1.2.1` does not match.

```
OID = .1.3.6.1.4.1.18997
Mask = .1.1.1.1.1.1.1
```

```
OID = .1.3.6.1.2.1
Mask = .1.1.0.1.0.1
```

To add a pairing of an OID subtree value and a mask value to a new or existing view, use the following format:

```
ibrix_snmpview -a -v VIEWNAME [-t {include|exclude}] -o OID_SUBTREE [-m MASK_BITS]
```

The subtree is added in the named view. For example, to add the X9000 Software private MIB to the view named `hp`, enter:

```
ibrix_snmpview -a -v hp -o .1.3.6.1.4.1.18997 -m .1.1.1.1.1.1.1
```

## Configuring groups and users

A *group* defines the access control policy on managed objects for one or more users. All users must belong to a group. Groups and users exist only in SNMPv3. Groups are assigned a security level, which enforces use of authentication and privacy, and specific read and write views to identify which managed objects group members can read and write.

The command to create a group assigns its SNMPv3 security level, read and write views, and context name. A *context* is a collection of managed objects that can be accessed by an SNMP entity. A related option, `-m`, determines how the context is matched. The format follows:

```
ibrix_snmpgroup -c -g GROUPNAME [-s {noAuthNoPriv|authNoPriv|authPriv}]  
[-r READVIEW] [-w WRITEVIEW] [-x CONTEXT_NAME] [-m {exact|prefix}]
```

For example, to create the group `group2` to require authorization, no encryption, and read access to the `hp` view, enter:

```
ibrix_snmpgroup -c -g group2 -s authNoPriv -r hp
```

The format to create a user and add that user to a group follows:

```
ibrix_snmpuser -c -n USERNAME -g GROUPNAME [-j {MD5|SHA}]  
[-k AUTHORIZATION_PASSWORD] [-y {DES|AES}] [-z PRIVACY_PASSWORD]
```

Authentication and privacy settings are optional. An authentication password is required if the group has a security level of either `authNoPriv` or `authPriv`. The privacy password is required if the group has a security level of `authPriv`. If unspecified, MD5 is used as the authentication algorithm and DES as the privacy algorithm, with no passwords assigned.

For example, to create `user3`, add that user to `group2`, and specify an authorization password for authorization and no encryption, enter:

```
ibrix_snmpuser -c -n user3 -g group2 -k auth-passwd -s authNoPriv
```

## Deleting elements of the SNMP configuration

All of the SNMP commands employ the same syntax for delete operations, using the `-d` flag to indicate that the following object is to be deleted. The following command deletes a list of hosts that were trapsinks:

```
ibrix_snmptrap -d -h lab15-12.domain.com,lab15-13.domain.com,lab15-14.domain.com
```

There are two restrictions on SNMP object deletions:

- A view cannot be deleted if it is referenced by a group.
- A group cannot be deleted if it is referenced by a user.

## Listing SNMP configuration information

All of the SNMP commands employ the same syntax for list operations, using the `-l` flag. For example:

```
ibrix_snmpgroup -l
```

This command lists the defined group settings for all SNMP groups. Specifying an optional group name lists the defined settings for that group only.

---

# 6 Configuring system backups

## Backing up the management console configuration

The management console configuration is automatically backed up whenever the cluster configuration changes. The backup takes place on the node hosting the active management console (or on the Management Server, if a dedicated management console is configured).

The backup file is stored at `<ibrixhome>/tmp/fmbbackup.zip` on the machine where it was created.

In an agile configuration, the active management console notifies the passive management console when a new backup file is available. The passive management console then copies the file to `<ibrixhome>/tmp/fmbbackup.zip` on the node on which it is hosted. If a management console is in maintenance mode, it will also be notified when a new backup file is created, and will retrieve it from the active management console.

You can create an additional copy of the backup file at any time. Run the following command, which creates a `fmbbackup.zip` file in the `$IBRIXHOME/log` directory:

```
$IBRIXHOME/bin/db_backup.sh
```

Once each day, a `cron` job rotates the `$IBRIXHOME/log` directory into the `$IBRIXHOME/log/daily` subdirectory. The `cron` job also creates a new backup of the management console configuration in both `$IBRIXHOME/tmp` and `$IBRIXHOME/log`.

If you need to force a backup, use the following command:

```
<installdirectory>/bin/ibrix_fm -B
```

---

### ❗ IMPORTANT:

You will need the backup file to recover from server failures or to undo unwanted configuration changes. Whenever the cluster configuration changes, be sure to save a copy of `fmbbackup.zip` in a safe, remote location such as a node on another cluster.

---

## Using NDMP backup applications

The NDMP backup feature can be used to back up and recover entire X9000 Software file systems or portions of a file system. You can use any supported NDMP backup application to perform the backup and recovery operations. (In NDMP terminology, the backup application is referred to as a Data Management Application, or DMA.) The DMA is run on a management station separate from the cluster and communicates with the cluster's file serving nodes over a configurable socket port.

The NDMP backup feature supports the following:

- NDMP protocol versions 3 and 4
- Two-way NDMP operations

- Three-way NDMP operations between two X9000 systems

Each file serving node functions as an NDMP Server and runs the NDMP Server daemon (ndmpd) process. When you start a backup or restore operation on the DMA, you can specify the node and tape device to be used for the operation.

Following are considerations for configuring and using the NDMP feature:

- When configuring your system for NDMP operations, attach your tape devices to a SAN and then verify that the file serving nodes to be used for backup/restore operations can see the appropriate devices.
- When performing backup operations, take hardware snapshots of your file systems and then back up the snapshots.

## Configuring NDMP parameters on the cluster

Certain NDMP parameters must be configured to enable communications between the DMA and the NDMP Servers in the cluster. To configure the parameters on the management console GUI, select **Cluster Configuration** from the Navigator, and then select **NDMP Backup**. The NDMP Configuration Summary shows the default values for the parameters. Click **Modify** to configure the parameters for your cluster on the Configure NDMP dialog box. See the online help for a description of each field.

The screenshot shows the 'Configure NDMP' dialog box with the following fields and values:

- Enable NDMP Sessions: Yes
- Minimum Port Number: 1025
- Maximum Port Number: 65535
- Listener Port Number: 10000
- Username: ndmp
- Password: ndmp
- Log Level: 0
- TCP Window Size (Bytes): 163840
- Concurrent Sessions: 128
- DMA IP Addresses: (empty)

Buttons: Add, IP Address, Delete, OK, Cancel, Help.

(\*) Required Value

To configure NDMP parameters from the CLI, use the following command:

```
ibrix_ndmpconfig -c [-d IP1,IP2,IP3,...] [-m MINPORT] [-x MAXPORT] [-n LISTENPORT]
[-u USERNAME] [-p PASSWORD] [-e {0=disable,1=enable}] -v {0=10}} [-w BYTES]
[-z NUMSESSIONS]
```

## NDMP process management

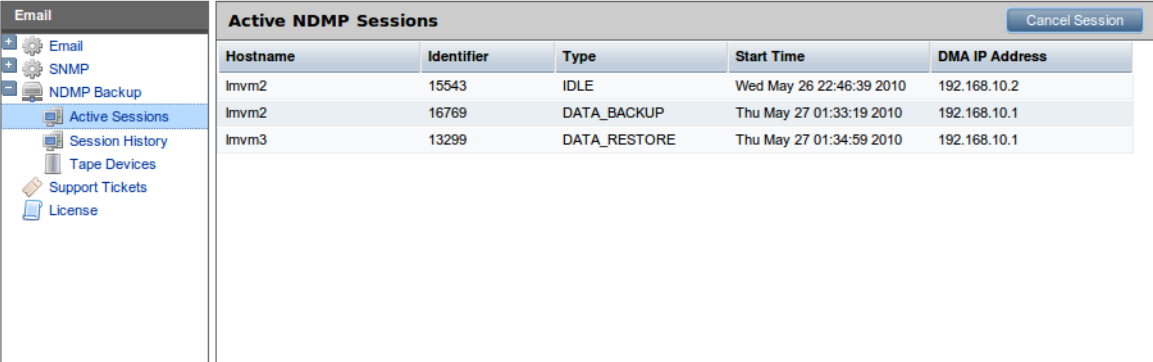
Normally all NDMP actions are controlled from the DMA. However, if the DMA cannot resolve a problem or you suspect that the DMA may have incorrect information about the NDMP environment, take the following actions from the X9000 Software management console GUI or CLI:

- Cancel one or more NDMP sessions on a file serving node. Canceling a session kills all spawned sessions processes and frees their resources if necessary.
- Reset the NDMP server on one or more file serving nodes. This step kills all spawned session processes, stops the ndmpd and session monitor daemons, frees all resources held by NDMP, and restarts the daemons.

### Viewing or canceling NDMP sessions

To view information about active NDMP sessions, select **Cluster Configuration** from the Navigator, and then select **NDMP Backup > Active Sessions**. For each session, the Active NDMP Sessions panel lists the host used for the session, the identifier generated by the backup application, the status of the session (backing up data, restoring data, or idle), the start time, and the IP address used by the DMA.

To cancel a session, select that session and click **Cancel Session**. Canceling a session kills all spawned sessions processes and frees their resources if necessary.



Hostname	Identifier	Type	Start Time	DMA IP Address
lmvm2	15543	IDLE	Wed May 26 22:46:39 2010	192.168.10.2
lmvm2	16769	DATA_BACKUP	Thu May 27 01:33:19 2010	192.168.10.1
lmvm3	13299	DATA_RESTORE	Thu May 27 01:34:59 2010	192.168.10.1

To see similar information for completed sessions, select **NDMP Backup > Session History**.

To view active sessions from the CLI, use the following command:

```
ibrix_ndmpsession -l
```

To view completed sessions, use the following command. The `-t` option restricts the history to sessions occurring on or before the specified date.

```
ibrix_ndmpsession -l -s [-t YYYY-MM-DD]
```

To cancel sessions on a specific file serving node, use the following command:

```
ibrix_ndmpsession -c SESSION1,SESSION2,SESSION3,... -h HOST
```

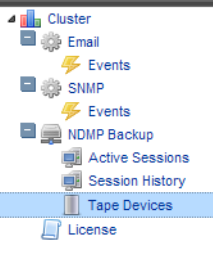
### Starting, stopping, or restarting an NDMP Server

When a file serving node is booted, the NDMP Server is started automatically. If necessary, you can use the following command to start, stop, or restart the NDMP Server on one or more file serving nodes:

```
ibrix_server -s -t ndmp -c { start | stop | restart} [-h SERVERNAMES]
```

## Viewing or rescanning tape and media changer devices

To view the tape and media changer devices currently configured for backups, select **Cluster Configuration** from the Navigator, and then select **NDMP Backup > Tape Devices**.

Cluster		Tape and Media Changer Devices			Rescan Devices
		Hostname	Device Type	Device ID	Device Node
 Cluster Email Events SNMP Events NDMP Backup Active Sessions Session History Tape Devices License		lmvm2	MediaChanger	HP:VLS:029AMWVPQ00	/dev/sg12
		lmvm2	TapeDrive	HP:Ultrium_3-SCSI:029AMWVPQ01	/dev/inst0
		lmvm2	TapeDrive	HP:Ultrium_3-SCSI:029AMWVPQ01	/dev/sg1
		lmvm2	TapeDrive	HP:Ultrium_3-SCSI:029AMWVPQ02	/dev/inst1
		lmvm2	TapeDrive	HP:Ultrium_3-SCSI:029AMWVPQ02	/dev/sg2
		lmvm2	TapeDrive	HP:Ultrium_3-SCSI:029AMWVPQ03	/dev/inst2
		lmvm2	TapeDrive	HP:Ultrium_3-SCSI:029AMWVPQ03	/dev/sg3
		lmvm2	TapeDrive	HP:Ultrium_3-SCSI:029AMWVPQ04	/dev/inst3

If you add a tape or media changer device to the SAN, click **Rescan Device** to update the list. If you remove a device and want to delete it from the list, you will need to reboot all of the servers to which the device is attached.

To view tape and media changer devices from the CLI, use the following command:

```
ibrix_tape -l
```

To rescan for devices, use the following command:

```
ibrix_tape -r
```

## NDMP events

An NDMP Server can generate three types of events: INFO, WARN, and ALERT. These events are displayed on the management console GUI and can be viewed with the `ibrix_event` command.

**INFO events.** These events specify when major NDMP operations start and finish, and also report progress. For example:

```
7012:Level 3 backup of /mnt/ibfs7 finished at Sat Nov 7 21:20:58 PST 2009
7013:Total Bytes = 38274665923, Average throughput = 236600391 bytes/sec.
```

**WARN events.** These events might indicate an issue with NDMP access, the environment, or NDMP operations. Be sure to review these events and take any necessary corrective actions. Following are some examples:

```
0000:Unauthorized NDMP Client 16.39.40.201 trying to connect
4002:User [joe] md5 mode login failed.
```

**ALERT events.** These alerts indicate that an NDMP action has failed. For example:

```
1102: Cannot start the session_monitor daemon, ndmpd exiting.
7009:Level 6 backup of /mnt/shares/accounts1 failed (writing eod header error).
8001:Restore Failed to read data stream signature.
```

You can configure the system to send email or SNMP notifications when these types of events occur.

---

# 7 Creating hostgroups for X9000 clients

A *hostgroup* is a named set of X9000 clients. Hostgroups provide a convenient way to centrally manage clients using the management console. You can put different sets of clients into hostgroups and then perform the following operations on all members of the group:

- Create and delete mountpoints
- Mount file systems
- Prefer a network interface
- Tune host parameters
- Set allocation policies

Hostgroups are optional. If you do not choose to set them up, you can mount file systems on clients and tune host settings and allocation policies on an individual level.

## How hostgroups work

In the simplest case, the hostgroups functionality allows you to perform an allowed operation on all X9000 clients by executing a management console command on the default `clients` hostgroup via either the CLI or the GUI. The `clients` hostgroup includes all X9000 clients configured in the cluster.



---

### NOTE:

The command intention is stored on the management console until the next time the clients contact the management console. (To force this contact, restart X9000 Software services on the clients, reboot them, or execute `ibrix_lwmount -a` or `ibrix_lwhost --a`.) When contacted, the management console informs the clients about commands that were executed on hostgroups to which they belong. The clients then use this information to perform the operation.

---

You can also use hostgroups to perform different operations on different sets of clients. To do this, you will need to create a *hostgroup tree* that includes the necessary hostgroups. You can then assign the clients manually, or the management console can automatically perform the assignment when you register an X9000 client, based on the client's cluster subnet. To use automatic assignment, you will need to create a domain rule that specifies the cluster subnet for the hostgroup.

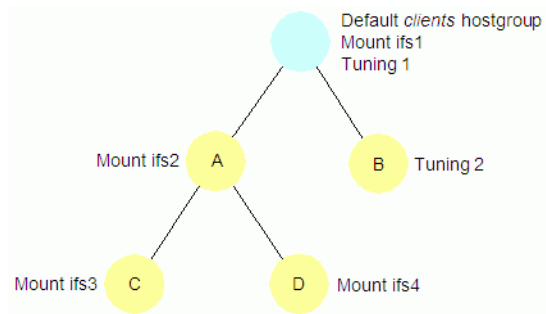
## Creating a hostgroup tree

The `clients` hostgroup is the root element of the hostgroup tree. Each hostgroup in a tree can have exactly one parent, and a parent can have multiple children, as shown in the following diagram). In a hostgroup tree, operations performed on lower-level nodes take precedence over operations performed on higher-level nodes. This means that you can effectively establish global client settings that you can override for specific clients.

For example, suppose that you want all clients to be able to mount file system `ifs1` and to implement a set of host tunings denoted as Tuning 1, but you want to override these global settings for certain

hostgroups. To do this, mount `ifs1` on the `clients` hostgroup, `ifs2` on hostgroup A, `ifs3` on hostgroup C, and `ifs4` on hostgroup D, in any order. Then, set Tuning 1 on the `clients` hostgroup and Tuning 2 on hostgroup B. The end result is that all clients in hostgroup B will mount `ifs1` and implement Tuning 2. The clients in hostgroup A will mount `ifs2` and implement Tuning 1. The clients in hostgroups C and D respectively, will mount `ifs3` and `ifs4` and implement Tuning 1.

The following diagram shows an example of global and local settings in a hostgroup tree.



To set up one level of hostgroups beneath the root, simply create the new hostgroups. You do not need to declare that the root node is the parent. To set up lower levels of hostgroups, declare a parent element for hostgroups.

Optionally, you can specify a domain rule for a hostgroup. Use only alphanumeric characters and the underscore character (`_`) in hostgroup names.

Do not use a host name as a group name.

To create a hostgroup tree using the CLI:

1. Create the first level of the tree and optionally declare a domain rule for it:

```
<installdirectory>/bin/ibrix_hostgroup -c -g GROUPNAME [-D DOMAIN]
```

2. Create all other levels by specifying a parent for the group and optionally a domain rule:

```
<installdirectory>/bin/ibrix_hostgroup -c -g GROUPNAME [-D DOMAIN] [-p PARENT]
```

## Adding an X9000 client to a hostgroup

You can add an X9000 client to a hostgroup or move a client to a different hostgroup. All clients belong to the default `clients` hostgroup.

To add or move a host to a hostgroup, use the `ibrix_hostgroup` command as follows:

```
<installdirectory>/bin/ibrix_hostgroup -m -g GROUP -h MEMBER
```

For example, to add the specified host to the `finance` group:

```
<installdirectory>/bin/ibrix_hostgroup -m -g finance -h cl01.hp.com
```

## Adding a domain rule to a hostgroup

To set up automatic hostgroup assignments, define a *domain rule* for hostgroups. A domain rule restricts hostgroup membership to clients on a particular cluster subnet. The management console uses the IP address that you specify for clients when you register them to perform a subnet match and sorts the clients into hostgroups based on the domain rules.

Setting domain rules on hostgroups provides a convenient way to centrally manage mounting, tuning, allocation policies, and preferred networks on different subnets of clients. A domain rule is a subnet



IP address that corresponds to a client network. Adding a domain rule to a hostgroup restricts its members to X9000 clients that are on the specified subnet. You can add a domain rule at any time.

To add a domain rule to a hostgroup, use the `ibrix_hostgroup` command as follows:

```
<installdirectory>/bin/ibrix_hostgroup -a -g GROUPNAME -D DOMAIN
```

For example, to add the domain rule 192.168 to the `finance` group:

```
<installdirectory>/bin/ibrix_hostgroup -a -g finance -D 192.168
```

## Viewing hostgroups

To view hostgroups, use the following command. You can view all hostgroups or a specific hostgroup.

```
<installdirectory>/bin/ibrix_hostgroup -l [-g GROUP]
```

## Deleting hostgroups

When you delete a hostgroup, its members are assigned to the parent of the deleted group.

To force the moved X9000 clients to implement the mounts, tunings, network interface preferences, and allocation policies that have been set on their new hostgroup, either restart X9000 Software services on the clients (see “Starting and stopping processes” in the system administration guide for your system) or execute the following commands locally:

- `ibrix_lwmount -a` to force the client to pick up mounts or allocation policies
- `ibrix_lwhost --a` to force the client to pick up host tunings

To delete a hostgroup using the CLI:

```
<installdirectory>/bin/ibrix_hostgroup -d -g GROUPNAME
```

## Other hostgroup operations

Additional hostgroup operations are described in the following locations:

- Creating or deleting a mountpoint, and mounting or unmounting a file system (see “Creating and mounting file systems” in the *HP StorageWorks X9000 File Serving Software File System User Guide*)
- Changing host tuning parameters (see [Tuning file serving nodes and X9000 clients](#), page 59)
- Preferring a network interface (see [Preferring network interfaces](#), page 65)
- Setting allocation policy (see “Using file allocation” in the *HP StorageWorks X9000 File Serving Software File System User Guide*)



# 8 Monitoring cluster operations

## Monitoring the status of file serving nodes

The dashboard on the management console GUI displays information about the operational status of file serving nodes, including CPU, I/O, and network performance information.

To view status from the CLI, use the `ibrix_server -l` command. This command provides CPU, I/O, and network performance information and indicates the operational state of the nodes, as shown in the following sample output:

```
<installdirectory>/bin/ibrix_server -l
```

SERVER_NAME	STATE	CPU (%)	NET_IO (MB/s)	DISK_IO (MB/s)	BACKUP	HA
node1	Up, HBAsDown	0	0.00	0.00		off
node2	Up, HBAsDown	0	0.00	0.00		off

File serving nodes can be in one of three operational states: Normal, Alert, or Error. These states are further broken down into categories that are mostly related to the failover status of the node. The following table describes the states.

State	Description
Normal	<b>Up:</b> Operational.
Alert	<b>Up-Alert:</b> Server has encountered a condition that has been logged. An event will appear in the Status tab of the management console GUI, and an email notification may be sent. <b>Up-InFailover:</b> Server is powered on and visible to the management console, and the management console is failing over the server's segments to a standby server. <b>Up-FailedOver:</b> Server is powered on and visible to the management console, and failover is complete.
Error	<b>Down-InFailover:</b> Server is powered down or inaccessible to the management console, and the management console is failing over the server's segments to a standby server. <b>Down-FailedOver:</b> Server is powered down or inaccessible to the management console, and failover is complete. <b>Down:</b> Server is powered down or inaccessible to the management console, and no standby server is providing access to the server's segments.

The STATE field also reports the status of monitored NICs and HBAs. If you have multiple HBAs and NICs and some of them are down, the state will be reported as HBAsDown or NicsDown.

## Monitoring cluster events

X9000 Software events are assigned to one of the following categories, based on the level of severity:

- **Alerts.** A disruptive event that can result in loss of access to file system data. For example, a segment is unavailable or a server is unreachable.
- **Warnings.** A potentially disruptive condition where file system access is not lost, but if the situation is not addressed, it can escalate to an alert condition. Some examples are reaching a very high server CPU utilization or nearing a quota limit.
- **Information.** An event that changes the cluster (such as creating a segment or mounting a file system) but occurs under normal or nonthreatening conditions.

Events are written to an events table in the configuration database as they are generated. To maintain the size of the file, HP recommends that you periodically remove the oldest events. See [“Removing events from the events database table”](#) on page 52 for more information.

You can set up event notifications through email (see [“Setting up email notification of cluster events”](#) on page 37) or SNMP traps (see [“Setting up SNMP notifications”](#) on page 38).

## Viewing events

The dashboard on the management console GUI specifies the number of events that have occurred in the last 24 hours. Click **Events** in the GUI Navigator to view a report of the events. You can also view events that have been reported for specific file systems or servers.

To view events from the CLI, use the following commands:

- View events by type:

```
<installdirectory>/bin/ibrix_event -q [-e ALERT|WARN|INFO]
```

- View generated events on a last-in, first-out basis:

```
<installdirectory>/bin/ibrix_event -l
```

- View a designated number of events. The command displays the 100 most recent messages by default. Use the `-n EVENTS_COUNT` option to increase or decrease the number of events displayed.

```
<installdirectory>/bin/ibrix_event -l [-n EVENTS_COUNT]
```

The following command displays the 25 most recent events:

```
<installdirectory>/bin/ibrix_event -l -n 25
```

## Removing events from the events database table

The `ibrix_event -p` command removes events from the events table, starting with the oldest events. The default is to remove the oldest seven days of events. To change the number of days, include the `-o DAYS_COUNT` option.

```
<installdirectory>/bin/ibrix_event -p [-o DAYS_COUNT]
```

## Monitoring cluster health

To monitor the functional health of file serving nodes and X9000 clients, execute the `ibrix_health` command. This command checks host performance in several functional areas and provides either a summary or a detailed report of the results.

## Health checks

The `ibrix_health` command runs these health checks on file serving nodes:

- Pings remote file serving nodes that share a network with the test hosts. Remote servers that are pingable might not be connected to a test host because of a Linux or X9000 Software issue. Remote servers that are not pingable might be down or have a network problem.
- If test hosts are assigned to be network interface monitors, pings their monitored interfaces to assess the health of the connection. (For information on network interface monitoring, see “[Using network interface monitoring](#)” on page 31.)
- Determines whether specified hosts can read their physical volumes.

The `ibrix_health` command runs this health check on both file serving nodes and X9000 clients:

- Determines whether information maps on the tested hosts are consistent with the configuration database.

If you include the `-b` option, the command also checks the health of standby servers (if configured).

## Health check reports

The summary report provides an overall health check result for all tested file serving nodes and X9000 clients, followed by individual results. If you include the `-b` option, the standby servers for all tested file serving nodes are included when the overall result is determined. The results will be one of the following:

- **Passed.** All tested hosts and standby servers passed every health check.
- **Failed.** One or more tested hosts failed a health check. The health status of standby servers is not included when this result is calculated.
- **Warning.** A suboptimal condition that might require your attention was found on one or more tested hosts or standby servers.

The detailed report consists of the summary report and the following additional data:

- Summary of the test results
- Host information such as operational state, performance data, and version data
- Nondefault host tunings
- Results of the health checks

By default, the Result Information field in a detailed report provides data only for health checks that received a Failed or a Warning result. Optionally, you can expand a detailed report to provide data about checks that received a Passed result, as well as details about the file system and segments.

## Viewing a summary health report

To view a summary health report, use the `ibrix_health -l` command:

```
<installdirectory>/bin/ibrix_health -l [-h HOSTLIST] [-f] [-b]
```

By default, the command reports on all hosts. To view specific hosts, include the `-h HOSTLIST` argument. To view results only for hosts that failed the check, include the `-f` argument. To include standby servers in the health check, include the `-b` argument.

For example, to view a summary report for node `i080` and client `lab13-116`:

```
<installdirectory>/bin/ibrix_health -l -h i080,lab13-116
```

Sample output follows:

```
PASSED
----- Host Summary Results -----
Host      Result  Type    State Last Update
=====  =====  =====  =====

```

```
i080      PASSED  Server  Up    Mon Apr 09 16:45:03 EDT 2007
lab13-116 PASSED  Client  Up    Mon Apr 09 16:07:22 EDT 2007
```

## Viewing a detailed health report

To view a detailed health report, use the `ibrix_health -i` command:

```
<installdirectory>/bin/ibrix_health -i -h HOSTLIST [-f] [-s] [-v]
```

The `-f` option displays results only for hosts that failed the check. The `-s` option includes information about the file system and its segments. The `-v` option includes details about checks that received a Passed or Warning result.

The following example shows a detailed health report for file serving node `lab13-116`:

```
<installdirectory>/bin/ibrix_health -i -h lab13-116
Overall Health Checker Results - PASSED
=====
Host Summary Results
=====
Host      Result  Type   State      Last Update
-----
lab15-62  PASSED  Server Up, HBAsDown Mon Oct 19 14:24:34 EDT 2009

lab15-62 Report
=====
Overall Result
=====
Result  Type   State      Module  Up time   Last Update
Network Thread Protocol
-----
PASSED  Server Up, HBAsDown Loaded  3267210.0 Mon Oct 19 14:24:34 EDT 2009
99.126.39.72 16      true

CPU Information
=====
Cpu(System,User,Util,Nice)  Load(1,3,15 min)  Network(Bps)  Disk(Bps)
-----
0, 1, 1, 0                  0.73, 0.17, 0.12   1301          9728

Memory Information
=====
Mem Total  Mem Free  Buffers(KB)  Cached(KB)  Swap Total(KB)  Swap Free(KB)
-----
1944532    1841548    688          34616       1028152         1028048

Version/OS Information
=====
Fs Version      IAD Version  OS           OS Version
-----
5.3.468(internal) 5.3.446     GNU/Linux   Red Hat Enterprise Linux Server release 5.2
(Tikanga) 2.6.18-92.el5 i386         i686

Remote Hosts
=====
Host      Type   Network      Protocol  Connection State
-----
lab15-61  Server 99.126.39.71 true      S_SET S_READY S_SENDHB
lab15-62  Server 99.126.39.72 true      S_NEW
```

```

Check Results
=====
Check : lab15-62 can ping remote segment server hosts
=====
Check Description          Result  Result Information
-----
Remote server lab15-61 pingable  PASSED

Check : Physical volumes are readable
=====
Check Description          Result  Result
Information
-----
Physical volume 0ownQk-vYcm-RziC-OwRU-qStr-C6d5-ESrMIf readable PASSED /dev/sde
Physical volume 1MY7Gk-zb7U-HnnA-D24H-Nxhg-WPmX-ZfUvMb readable PASSED /dev/sdc
Physical volume 7DRzC8-ucwo-p3D2-c89r-nwZD-Elju-6lVMw9 readable PASSED /dev/sda
Physical volume YipmIK-9WFE-tDpV-srtY-PoN7-9m23-r3Z9Gm readable PASSED /dev/sdb
Physical volume ansHXO-0zAL-K058-eEnZ-36ov-Pku2-Bz4Wks readable PASSED /dev/sdi
Physical volume oGt3qi-ybeC-E42f-vLg0-1GIF-My3H-3QhN0n readable PASSED /dev/sdj
Physical volume wzXSW3-2pxY-1ayt-2lkG-4yIH-fMez-QHfbgg readable PASSED /dev/sdd

Check : Iad and Fusion Manager consistent
=====
Check Description
-----
lab15-61 engine uuid matches on Iad and Fusion Manager
lab15-61 IP address matches on Iad and Fusion Manager
lab15-61 network protocol matches on Iad and Fusion Manager
lab15-61 engine connection state on Iad is up
lab15-62 engine uuid matches on Iad and Fusion Manager
lab15-62 IP address matches on Iad and Fusion Manager
lab15-62 network protocol matches on Iad and Fusion Manager
lab15-62 engine connection state on Iad is up
ifs2 file system uuid matches on Iad and Fusion Manager
ifs2 file system generation matches on Iad and Fusion Manager
ifs2 file system number segments matches on Iad and Fusion Manager
ifs2 file system mounted state matches on Iad and Fusion Manager
Segment owner for segment 1 filesystem ifs2 matches on Iad and Fusion Manager
Segment owner for segment 2 filesystem ifs2 matches on Iad and Fusion Manager
ifs1 file system uuid matches on Iad and Fusion Manager
ifs1 file system generation matches on Iad and Fusion Manager
ifs1 file system number segments matches on Iad and Fusion Manager
ifs1 file system mounted state matches on Iad and Fusion Manager
Segment owner for segment 1 filesystem ifs1 matches on Iad and Fusion Manager
Superblock owner for segment 1 of filesystem ifs2 on lab15-62 matches on Iad and
Fusion Manager PASSED
Superblock owner for segment 2 of filesystem ifs2 on lab15-62 matches on Iad and
Fusion Manager PASSED
Superblock owner for segment 1 of filesystem ifs1 on lab15-62 matches on Iad and
Fusion Manager PASSED

```

## Viewing logs

Logs are provided for the management console, file serving nodes, and X9000 clients. Contact HP Support for assistance in interpreting log files. You might be asked to tar the logs and email them to HP.

# Viewing operating statistics for file serving nodes

Periodically, the file serving nodes report the following statistics to the management console:

- **Summary.** General operational statistics including CPU usage, disk throughput, network throughput, and operational state. For information about the operational states, see [Monitoring the status of file serving nodes](#), page 51.
- **IO.** Aggregate statistics about reads and writes.
- **Network.** Aggregate statistics about network inputs and outputs.
- **Memory.** Statistics about available total, free, and swap memory.
- **CPU.** Statistics about processor and CPU activity.
- **NFS.** Statistics about NFS client and server activity.

The management console GUI displays most of these statistics on the dashboard. See “[Using the GUI](#)” on page 14 for more information.

To view the statistics from the CLI, use the following command:

```
<installdirectory>/bin/ibrix_stats -l [-s] [-c] [-m] [-i] [-n] [-f] [-h HOSTLIST]
```

Use the options to view only certain statistics or to view statistics for specific file serving nodes:

- s Summary statistics
- c CPU statistics
- m Memory statistics
- i I/O statistics
- n Network statistics
- f NFS statistics
- h The file serving nodes to be included in the report

Sample output follows:

```
-----Summary-----
HOST          Status  CPU  Disk(MB/s)  Net(MB/s)
lab12-10.hp.com  Up      0    22528      616
-----IO-----
HOST Read(MB/s)    Read(IO/s)  Read(ms/op)  Write(MB/s)  Write(IO/s)  Write(ms/op)
lab12-10.hp.com  22528      2           5           0           0.00
-----Net-----
HOST          In(MB/s)  In(IO/s)  Out(MB/s)  Out(IO/s)
lab12-10.hp.com  261      3        355      2
-----Mem-----
HOST          MemTotal(MB)  MemFree(MB)  SwapTotal(MB)  SwapFree(MB)
lab12-10.hp.com  1034616      703672      2031608      2031360
-----CPU-----
HOST          User  System  Nice  Idle  IoWait  Irq  SoftIrq
lab12-10.hp.com  0    0      0    0    97     1    0
-----NFS v3-----
HOST          Null  Getattr  Setattr  Lookup  Access  Readlink  Read  Write
lab12-10.hp.com  0    0      0      0      0      0      0    0

HOST          Create  Mkdir  Symlink  Mknod  Remove  Rmdir  Rename
lab12-10.hp.com  0      0      0      0      0      0      0

HOST          Link  Readdir  Readdirplus  Fsstat  Fsinfo  Pathconf  Commit
lab12-10.hp.com  0    0      0          0      0      0      0
```



---

# 9 Maintaining the system

## Shutting down the system

To shut down the system completely, first shut down the X9000 Software, and then power off the system hardware.

### Shutting down the X9000 Software

Use the following procedure to shut down the X9000 Software. Unless noted otherwise, run the commands from the dedicated Management Console or from the node hosting the active agile management console.

1. Disable HA for all file serving nodes:

```
ibrix_server -m -U
```

2. If your cluster has an agile management console configuration, place the passive management console into maintenance mode. Run the following command on the node hosting the passive management console:

```
ibrix_fm -m maintenance
```

3. Stop application services (CIFS, NFS, NDMP backup):

```
ibrix_server -s -t { cifs | nfs | ndmp } -c stop [-h SERVERLIST]
```

4. Unmount all file systems:

```
ibrix_umount -f <fs_name>
```

To unmount file systems from the management console GUI, select **Filesystems >unmount**.

5. Unmount all file systems from X9000 clients.

- On Linux X9000 clients, run the following command:

```
ibrix_lwumount -f <fs_name>
```

- On Windows X9000 clients, stop all applications accessing the file systems, and then use the client GUI to unmount the file systems (for example, I: DRIVE). Next, go to **Services** and stop the fusion service.

6. Verify that all file systems are unmounted:

```
ibrix_fs -l
```

7. Shut down file serving nodes other than the node hosting the active agile management console:

```
shutdown -t now "now"
```

8. Shut down the dedicated management console or the node hosting the active agile management console:

```
shutdown -t now "now"
```

## Powering off the hardware

Power off the file serving nodes in any order. The step completely shuts down the cluster.

## Starting the system

To start the system, first power on the file serving nodes, and then start the X900 Software.

### Starting the X9000 Software

To start the X9000 Software, complete the following steps:

1. Power on the dedicated Management Console or the node hosting the active agile management console.
2. Power on the file serving nodes (\*root segment = segment 1; power on owner first, if possible).
3. Monitor the nodes on the management console and wait for them all to report UP in the output from the following command:

```
ibrix_server -l
```

4. Mount file systems and verify their content. Run the following command on the Management Console or file serving node hosting the active agile management console:

```
ibrix_mount -f fs_name -m <mountpoint>
```

On Linux X9000 clients, run the following command:

```
ibrix_lwmount -f fsname -m <mountpoint>
```

5. Enable HA on the file serving nodes. Run the following command on the Management Console or file serving node hosting the active agile management console:

```
ibrix_server -m
```

6. On the node hosting the passive agile management console, move the console back to passive mode:

```
ibrix_fm -m passive
```

The X9000 Software is now available, and you can now access your file systems.

## Powering file serving nodes on or off

When file serving nodes are connected to properly configured power sources, the nodes can be powered on or off or can be reset remotely. To prevent interruption of service, set up standbys for the nodes (see “[Identifying standbys for file serving nodes](#)” on page 27), and then manually fail them over before powering them off (see “[Manually failing over a file serving node](#)” on page 30). Remotely powering off a file serving node does not trigger failover.

To power on, power off, or reset a file serving node, use the following command:

```
<installdirectory>/bin/ibrix_server -P {on|reset|off} -h HOSTNAME
```

## Starting and stopping processes

You can start, stop, and restart processes and can display status for the processes that perform internal X9000 Software functions. The following commands also control the operation of PostgreSQL on the machine. The PostgreSQL service is available at `/usr/local/ibrix/init/`.

To start and stop processes and view process status on the management console, use the following command:

```
/etc/init.d/ibrix_fusionmanager [start | stop | restart | status]
```

To start and stop processes and view process status on a file serving node, use the following command. In certain situations, a follow-up action is required after stopping, starting, or restarting a file serving node.

```
/etc/init.d/ibrix_server [start | stop | restart | status]
```

To start and stop processes and view process status on an X9000 client, use the following command:

```
/etc/init.d/ibrix_client [start | stop | restart | status]
```

## Tuning file serving nodes and X9000 clients

The default host tuning settings are adequate for most cluster environments. However, HP Support may recommend that you change certain file serving node or X9000 client tuning settings to improve performance.

Host tuning changes are executed immediately for file serving nodes. For X9000 clients, a tuning intention is stored in the management console. When X9000 Software services start on a client, the client queries the management console for the host tunings that it should use and then implements them. If X9000 Software services are already running on a client, you can force the client to query the management console by executing `ibrix_client` or `ibrix_lwhost --a` on the client, or by rebooting the client.

You can locally override host tunings that have been set on clients by executing the `ibrix_lwhost` command.

All management console commands for tuning hosts include the `-h HOSTLIST` option, which supplies one or more hostgroups. Setting host tunings on a hostgroup is a convenient way to tune a set of clients all at once. To set the same host tunings on all clients, specify the `clients` hostgroup.

---

### △ CAUTION:

Changing host tuning settings will alter file system performance. Contact HP Support before changing host tuning settings.

---

Use the `ibrix_host_tune` command to list or change host tuning settings:

- To list default values and valid ranges for all permitted host tunings:

```
<installdirectory>/bin/ibrix_host_tune -L
```

- To tune host parameters on nodes or hostgroups:

```
<installdirectory>/bin/ibrix_host_tune -S {-h HOSTLIST|-g GROUPLIST} -o OPTIONLIST
```

Contact HP Support to obtain the values for `OPTIONLIST`. List the options as `option=value` pairs, separated by commas. To set host tunings on all clients, include the `-g clients` option.

- To reset host parameters to their default values on nodes or hostgroups:

```
<installdirectory>/bin/ibrix_host_tune -U {-h HOSTLIST|-g GROUPLIST} [-n OPTIONS]
```

To reset all options on all file serving nodes, hostgroups, and X9000 clients, omit the `-h HOSTLIST` and `-n OPTIONS` options. To reset host tunings on all clients, include the `-g clients` option.

The values that are restored depend on the values specified for the `-h HOSTLIST` command:

- **File serving nodes.** The default file serving node host tunings are restored.
  - **X9000 clients.** The host tunings that are in effect for the default `clients` hostgroup are restored.
  - **Hostgroups.** The host tunings that are in effect for the parent of the specified hostgroups are restored.
- To list host tuning settings on file serving nodes, X9000 clients, and hostgroups, use the following command. Omit the `-h` argument to see tunings for all hosts. Omit the `-n` argument to see all tunings.

```
<installdirectory>/bin/ibrix_host_tune -l [-h HOSTLIST] [-n OPTIONS]
```

- To set the communications protocol on nodes and hostgroups, use the following command. To set the protocol on all X9000 clients, include the `-g clients` option.
- `<installdirectory>/bin/ibrix_host_tune -p {UDP|TCP} {-h HOSTLIST| -g GROUPLIST}`
- To set server threads on file serving nodes, hostgroups, and X9000 clients:

```
<installdirectory>/bin/ibrix_host_tune -t THREADCOUNT {-h HOSTLIST| -g GROUPLIST}
```

- To set admin threads on file serving nodes, hostgroups, and X9000 clients, use this command. To set admin threads on all X9000 clients, include the `-g clients` option.

```
<installdirectory>/bin/ibrix_host_tune -a THREADCOUNT {-h HOSTLIST| -g GROUPLIST}
```

## Tuning X9000 clients locally

**Linux clients.** Use the `ibrix_lwhost` command to tune host parameters. For example, to set the communications protocol:

```
<installdirectory>/bin/ibrix_lwhost --protocol -p {tcp|udp}
```

To list host tuning parameters that have been changed from their defaults:

```
<installdirectory>/bin/ibrix_lwhost --list
```

See the `ibrix_lwhost` command description in the *HP StorageWorks X9000 File Serving Software CLI Reference Guide* for other available options.

**Windows clients.** Click the **Tune Host** tab on the Windows X9000 client GUI. Tunable parameters include the NIC to prefer (the default is the cluster interface), the communications protocol (UDP or TCP), and the number of server threads to use. See the online help for the client if necessary.

## Migrating segments

To improve cluster performance, segment ownership can be transferred from one host to another through *segment migration*. Segment migration transfers segment ownership but it does not move segments from their physical locations in networked storage systems. Segment ownership is recorded on the physical segment itself, and the ownership data is part of the metadata that the management console distributes to file serving nodes and X9000 clients so that they can locate segments.

## Migrating specific segments

Use the following command to migrate ownership of the segments in *LVLIST* on file system *FSNAME* to a new host and update the source host:

```
<installdirectory>/bin/ibrix_fs -m -f FSNAME -s LVLIST -h HOSTNAME [-M] [-F] [-N]
```

To force the migration, include *-M*. To skip the source host update during the migration, include *-F*. To skip host health checks, include *-N*.

The following command migrates ownership of *ilv2* and *ilv3* in file system *ifs1* to *s1.hp.com*:

```
<installdirectory>/bin/ibrix_fs -m -f ifs1 -s ilv2,ilv3 -h s1.hp.com
```

## Migrating all segments from one host to another

Use the following command to migrate ownership of the segments in file system *FSNAME* that are owned by *HOSTNAME1* to *HOSTNAME2* and update the source host:

```
<installdirectory>/bin/ibrix_fs -m -f FSNAME -H HOSTNAME1,HOSTNAME2 [-M] [-F] [-N]
```

For example, to migrate ownership of all segments in file system *ifs1* that reside on *s1.hp.com* to *s2.hp.com*:

```
<installdirectory>/bin/ibrix_fs -m -f ifs1 -H s1.hp.com,s2.hp.com
```

## Removing storage from the cluster

Before removing storage that is used for an X9000 Software file system, you will need to evacuate the segments (or logical volumes) storing file system data. This procedure moves the data to other segments in the file system and is transparent to users or applications accessing the file system. When evacuating a segment, you should be aware of the following restrictions:

- Segment evacuation uses the file system rebalance operation. While the rebalance task is running, the system will prevent tasks from running on the same file system. Similarly, if another task is running on the file system, the rebalance task cannot be scheduled until the first task is complete.
- You cannot evacuate or remove the root segment (segment #1).
- The file system must be quiescent (no active I/O while a segment is being evacuated). Running this utility while the file system is active may result in data inconsistency or loss.
- If quotas are enabled on the affected file system, the quotas must be disabled during the rebalance operation.

To evacuate a segment, complete the following steps:

1. Identify the segment residing on the physical volume to be removed. Select **Storage** from the Navigator on the management console GUI. Note the file system and segment number on the affected physical volume. In the following example, physical volume *d1* is being retired. Segment 1 from file system *ifs1* uses that physical volume.

Storage						 Discover	 Remove
Physical Volume	Size(MB)	Volume Group	Logical Volume	File System	Seg Num		
d1	1500	ivg1	ilv1	ifs1	1		
d2	1500	ivg2	ilv2	ifs2	1		
d3	1500	ivg3	ilv3	ifs2	2		

2. Locate other segments on the file system that can accommodate the data being evacuated from the affected segment. Select the file system on the management console GUI and then select

**Segments** from the lower Navigator. If segments with adequate space are not available, add segments to the file system. In this example, the data from segment 1 will be evacuated to segments 2 and 3.

Segments <span>Assign to Tier Migrate Mark Bad</span>							
Segment	Logical Volum	Owner	Space (GB)	Used (%)	State	Tier	Type
1	ilv1	ib4.hp.com	1.31	34	OK		MIXED
2	ilv2	ib4.hp.com	1.31	14	OK		MIXED
3	ilv3	ib4.hp.com	1.31	4	OK		MIXED

3. If quotas are enabled on the file system, disable them:

```
ibrix_fs -q -D -f FSNAME
```

4. Evacuate the segment. Select the file system on the management console GUI and then select **Tasks > Rebalancer** from the lower Navigator. Click **Start** on the Task Summary page to open the Start Rebalancing dialog, and then open the Advanced tab. In the Source Segments column, select the segments to evacuate, and in the Destination Segments column, select the segments to receive the data. (If you do not select destination segments, the data is spread among the available segments.) Then click **Evacuate source segments**.

The image shows the 'Start Rebalancing' dialog box with the 'Advanced' tab selected. It has two columns: 'Source Segments' and 'Destination Segments'. In the 'Source Segments' column, segment 1 is selected with a checkmark. In the 'Destination Segments' column, segments 2 and 3 are selected with checkmarks. Below these columns is a checkbox labeled 'Evacuate source segments' which is also checked. At the bottom are 'Cancel' and 'OK' buttons.

The Task Summary window displays the progress of the rebalance operation and reports any errors. If you need to stop the operation, click **Stop**.

5. When the rebalance operation completes, remove the storage from the cluster:

```
ibrix_replicate -f FSNAME -b EVACUATED_SEGNUM
```

If you evacuated the root segment (segment 1 by default), include the `-F` option in the command.

The segment number associated with the storage is not reused.

6. If quotas were disabled on the file system, unmount the file system and then re-enable quotas using the following command:

```
ibrix_fs -q -E -f FSNAME
```

Then remount the file system.

To evacuate a segment using the CLI, use the `ibrix_rebalance -e` command, as described in the *HP StorageWorks X9000 File Serving Software CLI Reference Guide*.

# Maintaining networks

## Cluster and user network interfaces

X9000 Software supports the following logical network interfaces:

- **Cluster network interface.** This network interface carries management console traffic, traffic between file serving nodes, and traffic between file serving nodes and clients. A cluster can have only one cluster interface. For backup purposes, each file serving node and management console can have two cluster NICs.
- **User network interface.** This network interface carries traffic between file serving nodes and clients. Multiple user network interfaces are permitted.

The cluster network interface was created for you when your cluster was installed. For clusters with an agile management console configuration, a virtual interface is used for the cluster network interface. One or more user network interfaces may also have been created, depending on your site's requirements. You can add user network interfaces as necessary.

## Adding user network interfaces

Although the cluster network can carry traffic between file serving nodes and either NFS/CIFS or X9000 clients, you may want to create user network interfaces to carry this traffic. If your cluster must accommodate a mix of NFS/CIFS clients and X9000 clients, or if you need to segregate client traffic to different networks, you will need one or more user networks. In general, it is better to assign a user network for NFS/CIFS traffic because the cluster network cannot host the virtual interfaces (VIFs) required for NFS/CIFS failover. HP recommends that you use a Gigabit Ethernet port (or faster) for user networks.


When creating user network interfaces for file serving nodes, keep in mind that nodes needing to communicate for file system coverage or for failover must be on the same network interface. Also, nodes set up as a failover pair must be connected to the same network interface.

HP recommends that the default network be routed through the base User Network interface.

For a highly available cluster, HP recommends that you put NFS traffic on a dedicated user network and then set up automated failover for it (see [“Setting up automated failover”](#) on page 27). This method prevents interruptions to NFS traffic. If the cluster interface is used for NFS traffic and that interface fails on a file serving node, any NFS clients using the failed interface to access a mounted file system will lose contact with the file system because they have no knowledge of the cluster and cannot reroute requests to the standby for the node.

## Link aggregation and virtual interfaces

When creating a user network interface, you can use link aggregation to combine physical resources into a single VIF. VIFs allow you to provide many named paths within the larger physical resource, each of which can be managed and routed independently, as shown in the following diagram. See the network interface vendor documentation for any rules or restrictions required for link aggregation.



## Identifying a user network interface for a file serving node

To identify a user network interface for specific file serving nodes, use the `ibrix_nic` command. The interface name (*IFNAME*) can include only alphanumeric characters and underscores, such as `eth1`.

```
<installdirectory>/bin/ibrix_nic -a -n IFNAME -h HOSTLIST
```

If you are identifying a VIF, add the VIF suffix (`:nnnn`) to the physical interface name. For example, the following command identifies virtual interface `eth1:1` to physical network interface `eth1` on file serving nodes `s1.hp.com` and `s2.hp.com`:

```
<installdirectory>/bin/ibrix_nic -a -n eth1:1 -h s1.hp.com,s2.hp.com
```

When you identify a user network interface for a file serving node, the management console queries the node for its IP address, netmask, and MAC address and imports the values into the configuration database. You can modify these values later if necessary.

If you identify a VIF, the management console does not automatically query the node. If the VIF will be used only as a standby network interface in an automated failover setup, the management console will query the node the first time a network is failed over to the VIF. Otherwise, you must enter the VIF's IP address and netmask manually in the configuration database (see ["Setting network interface options in the configuration database"](#) on page 64). The management console does not require a MAC address for a VIF.

If you created a user network interface for X9000 client traffic, you will need to prefer the network for the X9000 clients that will use the network (see ["Preferring network interfaces"](#) on page 65).

## Setting network interface options in the configuration database

To make a VIF usable, execute the following command to specify the IP address and netmask for the VIF. You can also use this command to modify certain `ifconfig` options for a network.

```
<installdirectory>/bin/ibrix_nic -c -n IFNAME -h HOSTNAME [-I IPADDR] [-M NETMASK] [-B BCASTADDR] [-T MTU]
```

For example, to set netmask `255.255.0.0` and broadcast address `10.0.0.4` for interface `eth3` on file serving node `s4.hp.com`:

```
<installdirectory>/bin/ibrix_nic -c -n eth3 -h s4.hp.com -M 255.255.0.0 -B 10.0.0.4
```



## Preferring network interfaces

After creating a user network interface for file serving nodes or X9000 clients, you will need to *prefer* the interface for those nodes and clients. (It is not necessary to prefer a network interface for NFS or CIFS clients, because they can select the correct user network interface at mount time.)

When you prefer a user network interface for traffic from a source host to a destination host, traffic in the reverse direction remains defaulted to the cluster interface.

A network interface preference is executed immediately on file serving nodes. For X9000 clients, the preference intention is stored on the management console. When X9000 Software services start on a client, the client queries the management console for the network interface that has been preferred for it and then begins to use that interface. If the services are already running on X9000 clients when you prefer a network interface, you can force clients to query the management console by executing the command `ibrix_lwhost --a` on the client or by rebooting the client.

### Preferring a network interface for a file serving node or X9000 client

The first command prefers a network interface for a File Server Node; the second command prefers a network interface for a client.

```
<installdirectory>/bin/ibrix_server -n -h SRCHOST -A DESTHOST/IFNAME  
<installdirectory>/bin/ibrix_client -n -h SRCHOST -A DESTHOST/IFNAME
```

Execute this command once for each destination host that the file serving node or X9000 client should contact using the specified network interface (*IFNAME*). For example, to prefer network interface `eth3` for traffic from file serving node `s1.hp.com` to file serving node `s2.hp.com`:

```
<installdirectory>/bin/ibrix_server -n -h s1.hp.com -A s2.hp.com/eth3
```

### Preferring a network interface for a hostgroup

You can prefer an interface for multiple X9000 clients at one time by specifying a hostgroup. To prefer a user network interface for all X9000 clients, specify the `clients` hostgroup. After preferring a network interface for a hostgroup, you can locally override the preference on individual X9000 clients with the command `ibrix_lwhost`.

To prefer a network interface for a hostgroup, use the following command:

```
<installdirectory>/bin/ibrix_hostgroup -n -g HOSTGROUP -A DESTHOST/IFNAME
```

The destination host (*DESTHOST*) cannot be a hostgroup. For example, to prefer network interface `eth3` for traffic from all X9000 clients (the `clients` hostgroup) to file serving node `s2.hp.com`:

```
<installdirectory>/bin/ibrix_hostgroup -n -g clients -A s2.hp.com/eth3
```

## Unpreferring network interfaces

To return file serving nodes or X9000 clients to the cluster interface, unprefer their preferred network interface. The first command unprefers a network interface for a file serving node; the second command unprefers a network interface for a client.

```
<installdirectory>/bin/ibrix_server -n -h SRCHOST -D DESTHOST  
<installdirectory>/bin/ibrix_client -n -h SRCHOST -D DESTHOST
```

To unprefer a network interface for a hostgroup, use the following command:

```
<installdirectory>/bin/ibrix_client -n -g HOSTGROUP -A DESTHOST
```

## Making network changes

This section describes how to change IP addresses, change the cluster interface, manage routing table entries, and delete a network interface.

### Changing the IP address for a Linux X9000 client

After changing the IP address for a Linux X9000 client, you must update the X9000 Software configuration with the new information to ensure that the management console can communicate with the client. Use the following procedure:

1. Unmount the file system from the client.
2. Change the client's IP address.
3. Reboot the client or restart the network interface card.
4. Delete the old IP address from the configuration database:

```
<installdirectory>/bin/ibrix_client -d -h CLIENT
```

5. Re-register the client with the management console:

```
<installdirectory>/bin/register_client -p console_IPAddress -c clusterIF -n ClientName
```

6. Remount the file system on the client.

### Changing the IP address for the cluster interface on a dedicated management console

You must change the IP address for the cluster interface on both the file serving nodes and the management console.

1. If High Availability is enabled, disable it by executing `ibrix_server -m -U`.
2. Unmount the file system from all file serving nodes, and reboot.
3. On each file serving node, locally change the IP address of the cluster interface.
4. Change the IP address of the cluster interface for each file serving node:

```
<installdirectory>/bin/ibrix_nic -c -n IFNAME -h HOSTNAME [-I IPADDR]
```

5. Remount the file system.
6. Re-enable High Availability if necessary by executing `ibrix_server -m`.

### Changing the cluster interface

If you restructure your networks, you might need to change the cluster interface. The following rules apply when selecting a new cluster interface:

- The management console must be connected to all machines (including standby servers) that use the cluster network interface. Each file serving node and X9000 client must be connected to the management console by the same cluster network interface. A Gigabit (or faster) Ethernet port must be used for the cluster interface.
- X9000 clients must have network connectivity to the file serving nodes that manage their data and to the standbys for those servers. This traffic can use the cluster network interface or a user network interface.

To specify a new cluster interface for a cluster with a dedicated management console, use the following command:

```
<installdirectory>/bin/ibrix_nic -t -n IFNAME -h HOSTNAME
```

To specify a new virtual cluster interface for a cluster with an agile management console configuration, use the following command:

```
<installdirectory>/bin/ibrix_fm -c <VIF IP address> -d <VIF Device> -n <VIF Netmask>
-v cluster [-I <Local IP address_or_DNS hostname>]
```

## Managing routing table entries

X9000 Software supports one route for each network interface in the system routing table. Entering a new route for an interface overwrites the existing routing table entry for that interface.

### Adding a routing table entry

To add a routing table entry, use the following command:

```
<installdirectory>/bin/ibrix_nic -r -n IFNAME -h HOSTNAME -A -R ROUTE
```

The following command adds a route for virtual interface `eth2:232` on file serving node `s2.hp.com`, sending all traffic through gateway `gw.hp.com`:

```
<installdirectory>/bin/ibrix_nic -r -n eth2:232 -h s2.hp.com -A -R gw.hp.com
```

### Deleting a routing table entry

If you delete a routing table entry, it is not replaced with a default entry. A new replacement route must be added manually. To delete a route, use the following command:

```
<installdirectory>/bin/ibrix_nic -r -n IFNAME -h HOSTNAME -D
```

The following command deletes all routing table entries for virtual interface `eth0:1` on file serving node `s2.hp.com`:

```
<installdirectory>/bin/ibrix_nic -r -n eth0:1 -h s2.hp.com -D
```

## Deleting a network interface

Before deleting the interface used as the cluster interface on a file serving node, you must assign a new interface as the cluster interface. See [“Changing the cluster interface”](#) on page 66.

To delete a network interface, use the following command:

```
<installdirectory>/bin/ibrix_nic -d -n IFNAME -h HOSTLIST
```

The following command deletes interface `eth3` from file serving nodes `s1.hp.com` and `s2.hp.com`:

```
<installdirectory>/bin/ibrix_nic -d -n eth3 -h s1.hp.com,s2.hp.com
```

## Viewing network interface information

Executing the `ibrix_nic` command with no arguments lists all interfaces on all file serving nodes. Include the `-h` option to list interfaces on specific hosts.

```
<installdirectory>/bin/ibrix_nic -l -h HOSTLIST
```

The following table describes the fields in the output.

Field	Description
BACKUP_HOST	File serving node for the standby network interface.
BACKUP-IF	Standby network interface.
HOST	File serving node. An asterisk (*) denotes the management console.
IFNAME	Network interface on this file serving node.
IP_ADDRESS	IP address of this NIC.
LINKMON	Whether monitoring is on for this NIC.
MAC_ADDR	MAC address of this NIC.
ROUTE	IP address in routing table used by this NIC.
STATE	Network interface state.
TYPE	Network type (cluster or user).

When `ibrix_nic` is used with the `-i` option, it reports detailed information about the interfaces. Use the `-h` option to limit the output to specific hosts. Use the `-n` option to view information for a specific interface.

```
ibrix_nic -i [-h HOSTLIST] [-n NAME]
```

---

# 10 Migrating to an agile management console configuration

The agile management console configuration provides one active management console and one passive management console installed on different nodes in the cluster. The migration procedure configures the current Management Server machine as a host for an agile management console and installs another instance of the agile management console on a file serving node.

You can continue to use the Management Server machine as an agile management console, or you can install the agile management console on a second file serving node and repurpose the Management Server machine. Following are some possibilities for repurposing the machine:

- For X9300 Gateway systems, if the Management Server machine is a DL380 server and has the same hardware, CPU, and memory as the other DL380 file serving nodes in the cluster, you can convert the machine to a file serving node (there must be an even number of file serving nodes in the cluster). If the machine is not the same model and configuration as the existing file serving nodes, it cannot be converted to a file serving node.
- The Management Server machine can be removed from the cluster and used as an X9000 client. (For information about installing an X9000 client, see the *HP X9000 File Serving Software User Guide*.)
- The Management Server machine can be removed the cluster and used for other purposes.

To perform the migration, the X9000 installation code must be available. As delivered, this code is provided in `/tmp/X93xx/ibrix`. If this directory no longer exists, download the installation code from the HP support website for your storage system.

---

## ① IMPORTANT:

The migration procedure can be used only on clusters running HP X9000 File Serving Software 5.4 or later.

---

## Backing up the configuration

Before starting the migration to the agile management console configuration, make a manual backup of the management console configuration:

```
ibrix_fm -B
```

The resulting backup archive is located at `/usr/local/ibrix/tmp/fmbbackup.zip`. Save a copy of this archive in a safe, remote location, in case recovery is needed.

## Performing the migration

Complete the following steps on the Management Server:

1. The agile management console uses a virtual interface (VIF) IP address to enable failover and prevent any interruptions to file serving nodes and X9000 clients. The existing cluster NIC IP address becomes the permanent VIF IP address. Identify an unused IP address to use as the Cluster NIC IP address for the currently running management console. Then reconfigure the cluster and user networks:
  - Edit the `/etc/sysconfig/network-scripts/ifcfg-bond0` file. Change the IP address to the new, unused IP address and also ensure that `ONBOOT=Yes`.
  - If you are using X9000 clients over the user `bond1` network, edit the `/etc/sysconfig/network-scripts/ifcfg-bond1` file. Change the IP address to another unused, reserved IP address.

Run one of the following commands:

```
/etc/init.d/network restart
service network restart
```

If you are not at the local terminal you might have to reconnect using the new IP address. The following example shows the syntax for creating the agile management console cluster VIF (`bond0:1`) and, if using X9000 clients, the agile management console user VIF (`bond1:0`).

```
ifconfig bond0:1 172.16.3.1
ifconfig bond1:0 10.30.83.1
```

2. Configure the agile management console:

```
ibrix_fm -c <cluster_IP_addr> -d <cluster_VIF_device> -n
<cluster_VIF_netmask> -v cluster -I <local_IP_addr>
```

For example:

```
[root@x109s1 ~]# ibrix_fm -c 172.16.3.1 -d bond0:1 -n 255.255.248.0 -v cluster
-I 172.16.3.100
Command succeeded!
```

The original cluster IP address is now configured to the newly created cluster VIF device (`bond0:1`).

3. If you are using X9000 clients and you created the interface `bond1:0` in step 1, now set up the user network VIF, specifying the user VIF IP address and VIF device used in step 1.



#### NOTE:

This step does not apply to CIFS/NFS clients. If you are not using X9000 clients, you can skip this step.

Set up the user network VIF:

```
ibrix_fm -c <user_VIF_IP> -d <user_VIF_device> -n <user_VIF_netmask>
-v user
```

For example:

```
[root@x109s1 ~]# ibrix_fm -c 10.30.83.1 -d bond1:0 -n 255.255.0.0 -v user
Command succeeded
```

4. Restart the Fusionmanager services:

```
/etc/init.d/ibrix_fusionmanager restart
```
5. If the Management Server is currently running X9000 Software 5.4 or 5.4.1, skip this step. If the Management Server is currently running X9000 Software 5.5, install the file serving node software on the management server:

```
./ibrixinit -ts -C <local_cluster_iface_device> -I
<cluster_VIF_IP_Addr> -F
```

For example:

```
./ibrixinit -ts -C eth4 -i 172.16.3.100 -F
```

6. Register the Management Server hostname and IP address with the now “active” agile management console:

```
ibrix_fm -R <server_hostname> -I <IP_addr> -a <active_cluster_IP>
```

For example:

```
ibrix_fm -R x109s1 -I 172.16.3.100 -a 172.16.3.1
```

7. Verify that the agile management console is active:

```
ibrix_fm -i
```

For example:

```
[root@x109s1 ~]# ibrix_fm -i
FusionServer: x109s1 (active, quorum is running)
=====
Command succeeded!
```

8. Verify that there is only one management console in this cluster:

```
ibrix_fm -f
```

For example:

```
[root@x109s1 ~]# ibrix_fm -f
NAME      IP ADDRESS
-----
X109s1    172.16.3.100
Command succeeded!
```

9. To provide high availability for the management console, install a passive agile management console on an existing file serving node. In the command, the `-F` option forces the overwrite of the `new_lvm2_uuid` file that was installed with the X9000 Software. Run the following command on the file serving node:

```
<install_code_directory>/ibrixinit -tm -C <local_cluster_interface_device>
-v <cluster_VIF_IP> -m <cluster_netmask> -d <cluster_VIF_device> -w 9009
-M passive -F
```

For example:

```
[root@x109s3 ibrix]# <install_code_directory>/ibrixinit -tm -C bond0 -v 172.16.3.1
-m 255.255.248.0 -d bond0:0 -V 10.30.83.1 -N 255.255.0.0 -D bond1:0 -w 9009
-M passive -F
```

10. Verify that both management consoles are in the cluster:

```
ibrix_fm -f
```

For example:

```
[root@x109s3 ibrix]# ibrix_fm -f
NAME      IP ADDRESS
-----
x109s1    172.16.3.100
x109s3    172.16.3.3
Command succeeded!
```

11. Verify that the newly installed management console is in passive mode:

```
ibrix_fm -i
```

For example:

```
[root@x109s3 ibrix]# ibrix_fm -i
FusionServer: x109s3 (passive, quorum is running)
=====
Command succeeded
```

## Removing the dedicated Management Server

This procedure removes the dedicated Management Server machine from the cluster and installs another instance of the agile management console on a second file serving node.

Complete the following steps:

1. On the Management Server machine, place the agile management console into maintenance mode:

```
ibrix_fm -m maintenance
```

2. Verify that the management console is in maintenance mode:

```
ibrix_fm -i
```

For example:

```
[root@x109s1 ibrix]# ibrix_fm -i
FusionServer: x109s1 (maintenance, quorum not started)
=====
Command succeeded!
```

Verify that the passive management console has become the active management console. Run the `ibrix_fm -i` command on the file serving node hosting the passive management console (x109s3 in this example). It may take up to two minutes for the passive management console to become active.

For example:

```
[root@x109s3 ibrix]# ibrix_fm -i
FusionServer: x109s3 (active, quorum is running)
=====
Command succeeded!
```

At this point, the `ibrix_fm -f` command reports two agile management consoles, the passive console on the Management Server (x109s1) and the active console on the file serving node (x109s3).

```
[root@x109s3 ibrix]# ibrix_fm -f
NAME      IP ADDRESS
-----
x109s1    172.16.3.100
x109s3    172.16.3.3
Command succeeded!
```

3. Uninstall the management console from the Management Server machine:

```
/ibrix/ibrixinit -tm -U
```

4. Verify that the uninstalled management console is no longer registered. Run the following command from the file serving node hosting the newly active management console:

```
ibrix_fm -f
```

The command should now report only the agile management console on the file serving node.

```
[root@x109s3 ibrix]# ibrix_fm -f
NAME      IP ADDRESS
-----
```



```
x109s3 172.16.3.3  
Command succeeded!
```

5. Remove the Management Server machine from the cluster database:

```
ibrix_server -d -h HOSTNAME
```

6. To provide high availability for the management console, install a passive agile management console on another file serving node. In the command, the `-F` option forces the overwrite of the `new_lvm2_uuid` file that was installed with the X9000 Software.

```
/ibrix/ibrixinit -tm -C <local_cluster_interface_device> -v  
<cluster_VIF_IP> -m <cluster_netmask> -d <cluster_VIF_device> -w 9009  
-M passive -F
```

For example:

```
[root@x109s5 ibrix]# ./ibrixinit -tm -C bond0 -v 172.16.3.1 -m 255.255.248.0  
-d bond0 -V 10.30.83.1 -N 255.255.0.0 -D bond1:0 -w 9009 -M passive -F
```

7. Verify that the newly installed agile management console is in passive mode:

```
ibrix_fm -i
```

For example:

```
[root@x109s5 ibrix]# ibrix_fm -i  
FusionServer: x109s3 (passive, quorum is running)  
=====
```

Command succeeded

You can now physically remove the Management Server machine from the cluster.



---

# 11 Upgrading the X9000 Software

This chapter describes how to upgrade to the latest X9000 File Serving Software release. The management console and all file serving nodes must be upgraded to the new release at the same time. X9000 Clients are supported for one version beyond their release. For example, an X9000 5.3.2 client can run with a 5.4 X9000 server, but not with a 5.5 X9000 server.

---

## ❗ IMPORTANT:

Do not start new remote replication jobs while a cluster upgrade is in progress. If replication jobs were running before the upgrade started, the jobs will continue to run without problems after the upgrade completes.

---

Upgrades can be run either online or offline:

- **Online upgrades.** This procedure upgrades the software while file systems remain mounted. Before upgrading a file serving node, you will need to fail the node over to its backup node, allowing file system access to continue. This procedure cannot be used for major upgrades, but is appropriate for minor and maintenance upgrades.
- **Offline upgrades.** This procedure requires that file systems be unmounted on the node and that services be stopped. (Each file serving node may need to be rebooted if NFS or CIFS causes the unmount operation to fail.) You can then perform the upgrade. Clients experience a short interruption to file system access while each file serving node is upgraded.

You can use an automatic or a manual procedure to perform an offline upgrade. Online upgrades must be performed manually.

## Automatic upgrades

The automated upgrade procedure is run as an offline upgrade. When each file serving node is upgraded, all file systems are unmounted from the node and services are stopped. Clients will experience a short interruption to file system access while the node is upgraded.

All file serving nodes and management consoles must be up when you perform the upgrade. If a node or management console is not up, the upgrade script will fail and you will need to use a manual upgrade procedure instead. To determine the status of your cluster nodes, check the dashboard on the GUI.

---

## 📝 NOTE:

If you are currently running the 5.4 release with a standard management console and want to convert to an agile management console configuration, see [Chapter 10](#). Migrate to the agile management console first, and then perform the upgrade.

---

To upgrade all nodes in the cluster automatically, complete the following steps:

1. Check the dashboard on the management console GUI to verify that all nodes are up.

2. On the current active management console, move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous X9000 installation on this node, the installer is in `/root/ibrix`.
3. On the current active management console, expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
4. Change to the installer directory on the active management console, if necessary. Run the following command:  

```
./auto_ibrixupgrade
```

The upgrade script performs all necessary upgrade steps on every server in the cluster and logs progress in the `upgrade.log` file. The log file is located in the installer directory.
5. Upgrade X9000 clients. See [Upgrading Linux X9000 clients](#), page 88 and [Upgrading Windows X9000 clients](#), page 89.
6. If you received a new license from HP, install it as described in the “Licensing” chapter in this guide.

## Manual upgrades

### Upgrade paths

There are two manual upgrade paths: a standard upgrade and an agile upgrade.

- The standard upgrade is used on clusters having a dedicated Management Server machine or blade running the management console software.
- The agile upgrade is used on clusters having an agile management console configuration, where the management console software is installed in an active/passive configuration on two cluster nodes.

To determine whether you have an agile management console configuration, run the `ibrix_fm -i` command. If the output reports the status as `quorum is not configured`, your cluster does not have an agile configuration.

Be sure to use the upgrade procedure corresponding to your management console configuration:

- For standard upgrades, use [Standard upgrade for clusters with a dedicated Management Server machine or blade](#), page 77.
- For agile upgrades, use [Agile upgrade for clusters with an agile management console configuration](#), page 81.



---

#### NOTE:

If you are using a dedicated Management Server and want to convert to an agile management console configuration, see [Chapter 10](#). Complete the migration first, and then perform the upgrade using the agile upgrade procedure.

---

## Online and offline upgrades

Online and offline upgrade procedures are available for both the standard and agile upgrades:

- **Online upgrades.** This procedure upgrades the software while file systems remain mounted. Before upgrading a file serving node, you will need to fail the node over to its backup node, allowing

file system access to continue. This procedure cannot be used for major upgrades, but is appropriate for minor and maintenance upgrades.

- **Offline upgrades.** This procedure requires that you first unmount file systems and stop services. (Each file serving node may need to be rebooted if NFS or CIFS causes the unmount operation to fail.) You can then perform the upgrade. Clients will experience a short interruption to file system access while each file serving node is upgraded.

## Standard upgrade for clusters with a dedicated Management Server machine or blade

Use these procedures if your cluster has a dedicated Management Server machine or blade hosting the management console software. The X9000 Software 5.4.x to 5.5 upgrade can be performed either online or offline. Future releases may require offline upgrades.



### NOTE:

Be sure to read all instructions before starting the upgrade procedure.

## Standard online upgrade

The management console must be upgraded first. You can then upgrade file serving nodes and X9000 Clients in any order.

### Upgrading the management console

Complete the following steps on the Management Server machine or blade:

1. Disable automated failover on all file serving nodes:  

```
<ibrixhome>/bin/ibrix_server -m -U
```
2. Verify that automated failover is off:  

```
<ibrixhome>/bin/ibrix_server -l
```

In the output, the HA column should display `off`.
3. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous X9000 installation on this node, the installer is in `/root/ibrix`.
4. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
5. Change to the installer directory if necessary and run the upgrade:  

```
./ibrixupgrade -f
```
6. Verify that the management console is operational:  

```
/etc/init.d/ibrix_fusionmanager status
```

The `status` command should report that the correct services are running. The output is similar to this:

```
Fusion Manager Daemon (pid 18748) running...
```
7. Check `/usr/local/ibrix/log/fusionserver.log` for errors.

## Upgrading file serving nodes

After the management console has been upgraded, complete the following steps on each file serving node:

1. From the management console, manually fail over the file serving node:  

```
<ibrixhome>/bin/ibrix_server -f -p -h HOSTNAME
```

The node reboots automatically.
  2. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous X9000 installation on this node, the installer is in `/root/ibrix`.
  3. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
  4. Change to the installer directory if necessary and execute the following command:  

```
./ibrixupgrade -f
```

The upgrade automatically stops services and restarts them when the process is complete.
  5. When the upgrade is complete, verify that the X9000 Software services are running on the node:  

```
/etc/init.d/ibrix_server status
```

The output is similar to the following. If the IAD service is not running on your system, contact HP Support.

```
IBRIX Filesystem Drivers loaded
ibrcud is running.. pid 23325
IBRIX IAD Server (pid 23368) running...
```
  6. Verify that the `ibrix` and `ipfs` services are running:  

```
lsmod|grep ibrix
ibrix 2323332 0 (unused)
lsmod|grep ipfs
ipfs1 102592 0 (unused)
```

If either `grep` command returns empty, contact HP Support.
  7. From the management console, verify that the new version of X9000 Software FS/IAS is installed on the file serving node:  

```
<ibrixhome>/bin/ibrix_version -l -S
```
  8. If the upgrade was successful, failback the file serving node:  

```
<ibrixhome>/bin/ibrix_server -f -U -h HOSTNAME
```
  9. Repeat steps 1 through 8 for each file serving node in the cluster.
- After all file serving nodes have been upgraded and failed back, complete the upgrade.

## Completing the upgrade

1. From the management console, turn automated failover back on:  

```
<ibrixhome>/bin/ibrix_server -m
```
2. Confirm that automated failover is enabled:  

```
<ibrixhome>/bin/ibrix_server -l
```

In the output, HA displays `on`.

3. Upgrade X9000 Clients:
  - For Linux clients, see [Upgrading Linux X9000 clients](#), page 88.
  - For Windows clients, see [Upgrading Windows X9000 clients](#), page 89.
4. Verify that all version indicators match for file serving nodes and X9000 Clients. Run the following command from the management console:
 

```
<ibrixhome>/bin/ibrix_version -l
```

If there is a version mismatch, run the `/ibrix/ibrixupgrade -f` script again on the affected node, and then recheck the versions. The installation is successful when all version indicators match. If you followed all instructions and the version indicators do not match, contact HP Support.
5. Propagate a new segment map for the cluster:
 

```
<ibrixhome>/bin/ibrix_dbck -I -f FSNAME
```
6. Verify the health of the cluster:
 

```
<ibrixhome>/bin/ibrix_health -l
```

The output should specify `Passed / on`.
7. If you received a new license from HP, install it as described in the “Licensing” chapter in this guide.

## Standard offline upgrade

This upgrade procedure is appropriate for major upgrades. The management console must be upgraded first. You can then upgrade file serving nodes or X9000 Clients in any order.

### Preparing for the upgrade

1. From the management console, disable automated failover on all file serving nodes:
 

```
<ibrixhome>/bin/ibrix_server -m -U
```
2. From the management console, verify that automated failover is off. In the output, the `HA` column should display `off`.
 

```
<ibrixhome>/bin/ibrix_server -l
```
3. Stop the NFS and SMB services on all file serving nodes to prevent NFS and CIFS clients from timing out:
 

```
<ibrixhome>/bin/ibrix_server -s -t cifs -c stop
```

```
<ibrixhome>/bin/ibrix_server -s -t nfs -c stop
```

Verify that all likewise services are down on all file serving nodes:

```
ps -ef | grep likewise
```

Use `kill -9` to kill any likewise services that are still running.
4. From the management console, unmount all X9000 file systems:
 

```
<ibrixhome>/bin/ibrix_umount -f <fsname>
```

### Upgrading the management console

Complete the following steps on the management console:

1. Force a backup of the configuration:
 

```
<ibrixhome>/bin/ibrix_fm -B
```

The output is stored at `/usr/local/ibrix/tmp/fmbbackup.zip`. Be sure to save this file in a location outside of the cluster.

2. Move the <installer\_dir>/ibrix directory used in the previous release installation to ibrix.old. For example, if you expanded the tarball in /root during the previous X9000 installation on this node, the installer is in /root/ibrix.
3. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named ibrix that contains the installer program. For example, if you expand the tarball in /root, the installer is in /root/ibrix.
4. Change to the installer directory if necessary and execute the following command:  

```
./ibrixupgrade -f
```
5. Verify that the management console started successfully:  

```
/etc/init.d/ibrix_fusionmanager status
```

The status command confirms whether the correct services are running. Output is similar to the following:

```
Fusion Manager Daemon (pid 18748) running...
```
6. Check /usr/local/ibrix/log/fusionserver.log for errors.

### Upgrading the file serving nodes

After the management console has been upgraded, complete the following steps on each file serving node:

1. Move the <installer\_dir>/ibrix directory used in the previous release installation to ibrix.old. For example, if you expanded the tarball in /root during the previous X9000 installation on this node, the installer is in /root/ibrix.
2. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named ibrix that contains the installer program. For example, if you expand the tarball in /root, the installer is in /root/ibrix.
3. Change to the installer directory if necessary and execute the following command:  

```
./ibrixupgrade -f
```

The upgrade automatically stops services and restarts them when the process completes.
4. When the upgrade is complete, verify that the X9000 Software services are running on the node:  

```
/etc/init.d/ibrix_server status
```

The output should be similar to the following example. If the IAD service is not running on your system, contact HP Support.

```
IBRIX Filesystem Drivers loaded
ibrcud is running.. pid 23325
IBRIX IAD Server (pid 23368) running...
```
5. Execute the following commands to verify that the ibrix and ipfs services are running:  

```
lsmod|grep ibrix
ibrix 2323332 0 (unused)
lsmod|grep ipfs
ipfs1 102592 0 (unused)
```

If either grep command returns empty, contact HP Support.
6. From the management console, verify that the new version of X9000 Software FS/IAS has been installed on the file serving nodes:  

```
<ibrixhome>/bin/ibrix_version -l -S
```



## Completing the upgrade

1. Remount all file systems:  
`<ibrixhome>/bin/ibrix_mount -f <fsname> -m </mountpoint>`
2. From the management console, turn automated failover back on:  
`<ibrixhome>/bin/ibrix_server -m`
3. Confirm that automated failover is enabled:  
`<ibrixhome>/bin/ibrix_server -l`  
In the output, HA displays on.
4. From the management console, perform a manual backup of the upgraded configuration:  
`<ibrixhome>/bin/ibrix_fm -B`
5. Upgrade X9000 Clients:
  - For Linux clients, see [Upgrading Linux X9000 clients](#), page 88.
  - For Windows clients, see [Upgrading Windows X9000 clients](#), page 89.
6. Verify that all version indicators match for file serving nodes and X9000 Clients. Run the following command from the management console:  
`<ibrixhome>/bin/ibrix_version -l`  
If there is a version mismatch, run the `/ibrix/ibrixupgrade -f` script again on the affected node, and then recheck the versions. The installation is successful when all version indicators match. If you followed all instructions and the version indicators do not match, contact HP Support.
7. Verify the health of the cluster:  
`<ibrixhome>/bin/ibrix_health -l`  
The output should show `Passed / on`.
8. If you received a new license from HP, install it as described in the “Licensing” chapter in this document.

## Agile upgrade for clusters with an agile management console configuration

Use these procedures if your cluster has an agile management console configuration. The X9000 Software 5.4.x to 5.5 upgrade can be performed either online or offline. Future releases may require offline upgrades.

If you are currently running the 5.4 release with a standard management console and want to convert to an agile management console configuration, see [Chapter 10](#). Migrate to the agile management console first, and then perform the upgrade.



---

### NOTE:

Be sure to read all instructions before starting the upgrade procedure.

---

## Agile online upgrade

Perform the agile online upgrade in the following order:

- File serving node hosting the active management console
- File serving node hosting the passive management console
- Remaining file serving nodes and X9000 clients

## Upgrading the file serving nodes hosting the management console

Complete the following steps:

1. On the node hosting the active management console, force a backup of the management console configuration:  

```
<ibrixhome>/bin/ibrix_fm -B
```

The output is stored at `/usr/local/ibrix/tmp/fmbbackup.zip`. Be sure to save this file in a location outside of the cluster.
2. On the active management console node, disable automated failover on all file serving nodes:  

```
<ibrixhome>/bin/ibrix_server -m -U
```
3. Verify that automated failover is off:  

```
<ibrixhome>/bin/ibrix_server -l
```

In the output, the HA column should display `off`.
4. On the node hosting the active management console, place the management console into maintenance mode. This step fails over the active management console role to the node currently hosting the passive agile management console.  

```
<ibrixhome>/bin/ibrix_fm -m maintenance
```
5. Wait approximately 60 seconds for the failover to complete, and then run the following command on the node that was the target for the failover:  

```
<ibrixhome>/bin/ibrix_fm -i
```

The command should report that the agile management console is now `Active` on this node.
6. From the node on which you failed over the active management console in step 4, change the status of the management console from `maintenance` to `passive`:  

```
<ibrixhome>/bin/ibrix_fm -m passive
```
7. On the node hosting the active management console, manually fail over the node now hosting the passive management console:  

```
<ibrixhome>/bin/ibrix_server -f -p -h HOSTNAME
```

Wait a few minutes for the node to reboot and then run the following command to verify that the failover was successful. The output should report `Up, FailedOver`.  

```
<ibrixhome>/bin/ibrix_server -l
```
8. On the node hosting the active management console, place the management console into maintenance mode:  

```
<ibrixhome>/bin/ibrix_fm -m maintenance
```

This step fails back the active management console role to the node currently hosting the passive agile management console (the node that originally was active).
9. Wait approximately 90 seconds for the failover to complete, and then run the following command on the node that was the target for the failover:  

```
<ibrixhome>/bin/ibrix_fm -i
```

The command should report that the agile management console is now `Active` on this node.
10. On the node with the active agile management console, move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous X9000 installation on this node, the installer is in `/root/ibrix`.

11. On the node with the active agile management console, expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
12. Change to the installer directory if necessary and run the upgrade:  

```
./ibrixupgrade -f
```

The installer upgrades both the management console software and the file serving node software on this node.
13. Verify the status of the management console:  

```
/etc/init.d/ibrix_fusionmanager status
```

The status command confirms whether the correct services are running. Output will be similar to the following:

```
Fusion Manager Daemon (pid 18748) running...
```

Also run the following command, which should report that the console is Active:

```
<ibrixhome>/bin/ibrix_fm -i
```
14. Check `/usr/local/ibrix/log/fusionserver.log` for errors.
15. If the upgrade was successful, failback the file serving node. Run the following command on the node with the active agile management console:  

```
<ibrixhome>/bin/ibrix_server -f -U -h HOSTNAME
```
16. From the node on which you failed back the active management console in step 8, change the status of the management console from maintenance to passive:  

```
<ibrixhome>/bin/ibrix_fm -m passive
```
17. If the node with the passive management console is also a file serving node, manually fail over the node from the active management console:  

```
<ibrixhome>/bin/ibrix_server -f -p -h HOSTNAME
```

Wait a few minutes for the node to reboot, and then run the following command to verify that the failover was successful. The output should report `Up, FailedOver`.

```
<ibrixhome>/bin/ibrix_server -l
```
18. On the node with the passive agile management console, move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous X9000 installation on this node, the installer is in `/root/ibrix`.
19. On the node hosting the passive agile management console, expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
20. Change to the installer directory if necessary and run the upgrade:  

```
./ibrixupgrade -f
```

The installer upgrades both the management console software and the file serving node software on the node.
21. Verify the status of the management console:  

```
/etc/init.d/ibrix_fusionmanager status
```

The status command confirms whether the correct services are running. Output will be similar to the following:

```
Fusion Manager Daemon (pid 18748) running...
```

Also run the following command, which should report that the console is passive:

```
<ibrixhome>/bin/ibrix_fm -i
```

22. Check `/usr/local/ibrix/log/fusionserver.log` for errors.
23. If the upgrade was successful, fail back the node. Run the following command on the node with the active agile management console:

```
<ibrixhome>/bin/ibrix_server -f -U -h HOSTNAME
```

24. Verify that the agile management console software and the file serving node software are now upgraded on the two nodes hosting the agile management console:

```
<ibrixhome>/bin/ibrix_version -l -S
```

Following is some sample output:

```
Fusion Manager version: 5.5.XXX
=====
Segment Servers
=====
HOST_NAME  FILE_SYSTEM  IAD/IAS  IAD/FS  OS  KERNEL_VERSION  ARCH
-----
ib50-86    5.5.205(X9000_5_5)  5.5.XXX  5.5.XXX  GNU/Linux  2.6.18-128.el5  x86_64
ib50-87    5.5.205(X9000_5_5)  5.5.XXX  5.5.XXX  GNU/Linux  2.6.18-128.el5  x86_64
```

You can now upgrade any remaining file serving nodes and X9000 clients.

## Upgrading remaining file serving nodes

Complete the following steps on each file serving node:

1. Manually fail over the file serving node:

```
<ibrixhome>/bin/ibrix_server -f -p -h HOSTNAME
```

The node will be rebooted automatically.

2. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ib-rix.old`. For example, if you expanded the tarball in `/root` during the previous X9000 installation on this node, the installer is in `/root/ibrix`.
3. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
4. Change to the installer directory if necessary and execute the following command:

```
./ibrixupgrade -f
```

The upgrade automatically stops services and restarts them when the process is complete.

5. When the upgrade is complete, verify that the X9000 Software services are running on the node:

```
/etc/init.d/ibrix_server status
```

The output will be similar to the following. If the IAD service is not running on your system, contact HP Support.

```
IBRIX Filesystem Drivers loaded
ibrcud is running.. pid 23325
IBRIX IAD Server (pid 23368) running...
```

6. Verify that the `ibrix` and `ipfs` services are running:

```
lsmod|grep ibrix
ibrix 2323332 0 (unused)
lsmod|grep ipfs
```

```
ipfs1 102592 0 (unused)
```

If either `grep` command returns empty, contact HP Support.

7. From the management console, verify that the new version of X9000 Software FS/IAS has been installed on the file serving node:

```
<ibrixhome>/bin/ibrix_version -l -S
```

8. If the upgrade was successful, failback the file serving node:

```
<ibrixhome>/bin/ibrix_server -f -U -h HOSTNAME
```

9. Repeat steps 1 through 8 for each remaining file serving node in the cluster.

After all file serving nodes have been upgraded and failed back, complete the upgrade.

### Completing the upgrade

1. From the node hosting the active management console, turn automated failover back on:

```
<ibrixhome>/bin/ibrix_server -m
```

2. Confirm that automated failover is enabled:

```
<ibrixhome>/bin/ibrix_server -l
```

In the output, the HA column should display `on`.

3. Upgrade X9000 Clients:

- For Linux clients, see [Upgrading Linux X9000 clients](#), page 88.
- For Windows clients, see [Upgrading Windows X9000 clients](#), page 89.

4. Verify that all version indicators match for file serving nodes and X9000 Clients. Run the following command from the active management console:

```
<ibrixhome>/bin/ibrix_version -l
```

If there is a version mismatch, run the `/ibrix/ibrixupgrade -f` script again on the affected node, and then recheck the versions. The installation is successful when all version indicators match. If you followed all instructions and the version indicators do not match, contact HP Support.

5. Propagate a new segment map for the cluster:

```
<ibrixhome>/bin/ibrix_dbck -I -f FSNAME
```

6. Verify the health of the cluster:

```
<ibrixhome>/bin/ibrix_health -l
```

The output should specify `Passed / on`.

7. If you received a new license from HP, install it as described in the “Licensing” chapter in this guide.

### Agile offline upgrade

This upgrade procedure is appropriate for major upgrades. Perform the agile offline upgrade in the following order:

- File serving node hosting the active management console
- File serving node hosting the passive management console
- Remaining file serving nodes and X9000 clients

**NOTE:**

To determine which node is hosting the active management console, run the following command:

```
<ibrixhome>/bin/ibrix_fm -i
```

## Preparing for the upgrade

1. On the active management console node, disable automated failover on all file serving nodes:  

```
<ibrixhome>/bin/ibrix_server -m -U
```
2. Verify that automated failover is off. In the output, the HA column should display `off`.  

```
<ibrixhome>/bin/ibrix_server -l
```
3. On the active management console node, stop the NFS and SMB services on all file serving nodes to prevent NFS and CIFS clients from timing out.  

```
<ibrixhome>/bin/ibrix_server -s -t cifs -c stop
```

```
<ibrixhome>/bin/ibrix_server -s -t nfs -c stop
```

Verify that all likewise services are down on all file serving nodes:

```
ps -ef | grep likewise
```

Use `kill -9` to kill any likewise services that are still running.
4. Unmount all X9000 Software file systems:  

```
<ibrixhome>/bin/ibrix_umount -f <fsname>
```

## Upgrading the file serving nodes hosting the management console

Complete the following steps:

1. On the node hosting the active management console, force a backup of the management console configuration:  

```
<ibrixhome>/bin/ibrix_fm -B
```

The output is stored at `/usr/local/ibrix/tmp/fmbbackup.zip`. Be sure to save this file in a location outside of the cluster.
2. On the node hosting the passive management console, place the management console into maintenance mode:  

```
<ibrixhome>/bin/ibrix_fm -m maintenance
```
3. On the active management console node, move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous X9000 installation on this node, the installer is in `/root/ibrix`.
4. On the active management console node, expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
5. Change to the installer directory if necessary and run the upgrade:  

```
./ibrixupgrade -f
```

The installer upgrades both the management console software and the file serving node software on this node.
6. Verify the status of the management console:

```
/etc/init.d/ibrix_fusionmanager status
```

The status command confirms whether the correct services are running. Output will be similar to the following:

```
Fusion Manager Daemon (pid 18748) running...
```

7. Check `/usr/local/ibrix/log/fusionserver.log` for errors.
8. Upgrade the remaining management console node. Move the `ibrix` directory used in the previous release to `ibrix.old`. Then expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
9. Change to the installer directory if necessary and run the upgrade:  

```
./ibrixupgrade -f
```

The installer upgrades both the management console software and the file serving node software on the node.
10. On the node that was just upgraded and has its management console in maintenance mode, move the management console back to passive mode:  

```
<ibrixhome>/bin/ibrix_fm -m passive
```

The node now resumes its normal backup operation for the active management console.

### Upgrading remaining file serving nodes

Complete the following steps on the remaining file serving nodes:

1. Move the `<installer_dir>/ibrix` directory used in the previous release installation to `ibrix.old`. For example, if you expanded the tarball in `/root` during the previous X9000 installation on this node, the installer is in `/root/ibrix`.
2. Expand the distribution tarball or mount the distribution DVD in a directory of your choice. Expanding the tarball creates a subdirectory named `ibrix` that contains the installer program. For example, if you expand the tarball in `/root`, the installer is in `/root/ibrix`.
3. Change to the installer directory if necessary and execute the following command:  

```
./ibrixupgrade -f
```

The upgrade automatically stops services and restarts them when the process is complete.
4. When the upgrade is complete, verify that the X9000 Software services are running on the node:  

```
/etc/init.d/ibrix_server status
```

The output should be similar to the following example. If the IAD service is not running on your system, contact HP Support.

```
IBRIX Filesystem Drivers loaded
ibrcud is running.. pid 23325
IBRIX IAD Server (pid 23368) running...
```
5. Execute the following commands to verify that the `ibrix` and `ipfs` services are running:  

```
lsmod|grep ibrix
ibrix 2323332 0 (unused)
lsmod|grep ipfs
ipfs1 102592 0 (unused)
```

If either `grep` command returns empty, contact HP Support.

6. From the active management console node, verify that the new version of X9000 Software FS/IAS is installed on the file serving nodes:

```
<ibrixhome>/bin/ibrix_version -l -S
```

### Completing the upgrade

1. Remount the X9000 Software file systems:  

```
<ibrixhome>/bin/ibrix_mount -f <fsname> -m </mountpoint>
```
2. From the node hosting the active management console, turn automated failover back on:  

```
<ibrixhome>/bin/ibrix_server -m
```
3. Confirm that automated failover is enabled:  

```
<ibrixhome>/bin/ibrix_server -l
```

In the output, HA should display on.
4. From the node hosting the active management console, perform a manual backup of the upgraded configuration:  

```
<ibrixhome>/bin/ibrix_fm -B
```
5. Upgrade X9000 Clients:
  - For Linux clients, see [Upgrading Linux X9000 clients](#), page 88.
  - For Windows clients, see [Upgrading Windows X9000 clients](#), page 89.
6. Verify that all version indicators match for file serving nodes and X9000 Client. Run the following command from the active management console:  

```
<ibrixhome>/bin/ibrix_version -l
```

If there is a version mismatch, run the `/ibrix/ibrixupgrade -f` script again on the affected node, and then recheck the versions. The installation is successful when all version indicators match. If you followed all instructions and the version indicators do not match, contact HP Support.
7. Verify the health of the cluster:  

```
<ibrixhome>/bin/ibrix_health -l
```

The output should show `Passed / on`.
8. If you received a new license from HP, install it as described in the “Licensing” chapter in this guide.

## Upgrading Linux X9000 clients

Be sure to upgrade the management console and file serving nodes before upgrading Linux X9000 clients. Complete the following steps on each client:

1. Expand the upgrade tarball or mount the upgrade DVD.
2. Run the upgrade script:  

```
./ibrixupgrade -f
```

The upgrade software automatically stops the necessary services and restarts them when the upgrade is complete.
3. Execute the following command to verify the client is running X9000 Software:  

```
/etc/init.d/ibrix_client status  
IBRIX Filesystem Drivers loaded  
IBRIX IAD Server (pid 3208) running...
```



The `IAD` service should be running, as shown in the sample output above. If it is not, contact HP Support.

## Upgrading Windows X9000 clients

Complete the following steps on each client:

1. Remove the old Windows X9000 client software using the **Add or Remove Programs** utility in the Control Panel.
2. Copy the Windows X9000 client MSI file for the upgrade to the machine.
3. Launch the Windows Installer and follow the instructions to complete the upgrade.
4. Check **Administrative Tools | Services** to verify that the X9000 Client service is started.
5. Launch the Windows X9000 client. On the **Active Directory Settings** tab, click **Update** to retrieve the current settings.



### NOTE:

If you are using Remote Desktop to perform an upgrade, you must log out and log back in to see the drive mounted.

## Troubleshooting upgrade issues

### Automatic upgrade fails

Check the `upgrade.log` file to determine the source of the failure. (The log file is located in the installer directory.) If it is not possible to perform the automatic upgrade, continue with the manual upgrade procedure.

### ibrixupgrade hangs

The installation can hang because the RPM database is corrupted. This is caused by inconsistencies in the Red Hat Package Manager.

Rebuild the RPM database using the following commands and then attempt the installation again. Note that `rm` is followed by a space and then two underscores, and `rpm` is followed by a space and then two dashes:

```
cd /var/lib/rpm
rm __*
rpm --rebuilddb
```

On the management console, `ibrixupgrade` may also hang if the NFS mount points are stale. In this case, clean up the mount points, reboot the management console, and run the upgrade procedure again.

### “Access denied” error on Windows Client

After completing the Windows Client upgrade, you may receive an `access denied` error when writing to a cluster drive. To recover write access, you will need to recover the Client's registration.

To do this, launch the X9000 Software Client User Interface on the Client. Go to the Registration Tab, enter the Management Console name, select **Recover Registration**, and then click **Register**. If you are prompted to overwrite the existing registration, select **yes** to complete the operation.

---

# 12 Licensing

This chapter describes how to view your current license terms and how to obtain and install new X9000 Software product license keys.



---

## NOTE:

For MSA2000 G2 licensing (for example, snapshots), see the MSA2000 G2 documentation.

---

## Viewing license terms

The X9000 Software license file is stored in the installation directory on the management console. To view the license from the management console GUI, select **Cluster Configuration** in the Navigator and then select **License**.

To view the license from the CLI, use the following command:

```
<installdirectory>/bin/ibrix_license -i
```

The output reports your current node count and capacity limit. In the output, Segment Server refers to file serving nodes.

## Retrieving a license key

When you purchased this product, you received a License Entitlement Certificate. You will need information from this certificate to retrieve and enter your license keys.

You can use any of the following methods to request a license key:

- Obtain a license key from <http://webware.hp.com>.
- Use AutoPass to retrieve and install permanent license keys. See [Using AutoPass to retrieve and install permanent license keys](#).
- Fax the Password Request Form that came with your License Entitlement Certificate. See the certificate for fax numbers in your area.
- Call or email the HP Password Center. See the certificate for telephone numbers in your area or email addresses.

## Using AutoPass to retrieve and install permanent license keys

The procedure must be run from a client with JRE 1.5 or later installed and with a desktop manager running (for example, a Linux-based system running X Windows). The `ssh` client must also be installed.

1. On the Linux-based system, run the following command to connect to the Management Console:

```
ssh -X root@<management_console_IP>
```

2. When prompted, enter the password for the management console.

3. Launch the AutoPass GUI:

```
/usr/local/ibrx/bin/fusion-license-manager
```

4. In the AutoPass GUI, go to **Tools**, select **Configure Proxy**, and configure your proxy settings.
5. Click **Retrieve/Install License > Key** and then retrieve and install your license key.

If the management console does not have an Internet connection, retrieve the license from a machine that does have a connection, deliver the file with the license to the management console machine, and then use the AutoPass GUI to import the license.

---

# 13 Upgrading firmware

## Upgradable firmware

The HP X9300 system includes several components with upgradable firmware. The following table lists these components and specifies whether they can be upgraded online and in a nondisruptive manner.

Component	Online and Nondisruptive?
DL380	Nondisruptive if done one server at a time
FC HBA (HP X9300 systems)	Yes, if done one server at a time
SAS HBA	Yes, if done one server at a time
OS image	Yes, if done one server at a time
RAID controller	Yes, if done one controller at a time
HDD	No

## Installing firmware upgrades

See the documentation provided with the firmware for installation instructions.



---

# 14 Troubleshooting

## Managing support tickets

A support ticket includes system and X9000 software information useful for analyzing performance issues and node terminations. A support ticket is created automatically if a file serving node terminates unexpectedly. You can also create a ticket manually if your cluster experiences issues that need to be investigated by HP Support.

The collected information is collated into a tar file and placed in the directory `/admin/platform/diag/support/tickets/` on the active management console. Send this tar file to HP Support for analysis. The name of the tar file is `ticket_<name>.tgz`. In the filename, `<name>` is a number, for example, `ticket_0002.tgz`. To view or delete a specific ticket, use the name assigned to the ticket.

The Support Ticket feature requires that one-way shared SSH keys be configured on all file serving nodes. For new systems, the keys were configured for you when the cluster was installed. If you upgraded from a release earlier than 5.4, you need to configure the keys. (See [“Configuring shared ssh keys”](#) on page 97.)

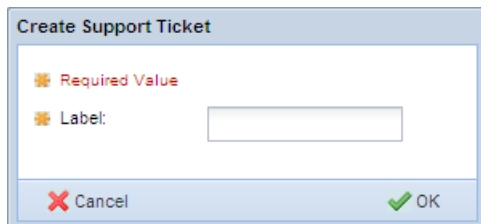


### NOTE:

When the cluster includes an agile management console configuration, the Support Ticket information shown on the management console GUI or CLI is in the context of the currently active management console. If the active management console fails over and the passive management console becomes active, the existing support ticket information does not move to the newly active management console. Support Ticket operations are always handled by the currently active management console and the final output of the operations is stored there.

## Creating, viewing, and deleting support tickets

To create a support ticket, select **Support Ticket** from the GUI Navigator, and then select **Create** from the Options list. On the Create Support Ticket dialog box, enter a label to help identify the ticket. The label is for your information only.



The image shows a 'Create Support Ticket' dialog box. It has a title bar with the text 'Create Support Ticket'. Inside the dialog, there is a label 'Required Value' in red text. Below it, there is a label 'Label:' followed by a text input field. At the bottom of the dialog, there are two buttons: 'Cancel' with a red 'X' icon and 'OK' with a green checkmark icon.

To create a support ticket from the CLI, use the following command:

```
<ibrixhome>/bin/ibrix_supportticket -c -L <Label>
```

To view a support ticket on the GUI, select **Support Tickets** from the Navigator. On the CLI, use the following command to view all support tickets:

```
<ibrixhome>/bin/ibrix_supportticket -l
```

To view details for a specific support ticket, use the following command:

```
<ibrixhome>/bin/ibrix_supportticket -v -n <Name>
```

When you no longer need a support ticket, you can delete it. From the GUI, select **Support Ticket** from the Navigator. Select the appropriate support ticket, select **Delete** from the Options menu, and confirm the operation.

To delete a support ticket from the CLI, use the following command:

```
<ibrixhome>/bin/ibrix_supportticket -d -n <Name>
```

## Support ticket states

Support tickets are in one of the following states:

Ticket State	Description
COLLECTING_LOGS	The data collection operation is collecting logs and command output.
COLLECTED_LOGS	The data collection operation has completed on all nodes in the cluster.
CREATING	The data collected from each node is being copied to the active management console.
CREATED	The ticket was created successfully. The data from each node is available in a tar file in the <code>/admin/platform/diag/support_tickets/</code> directory on the active management console.
PARTIALLY_CREATED	The ticket was created successfully. Certain nodes were unavailable at the time of copy; however, the data from the available nodes is available in a tar file in the directory <code>/admin/platform/diag/support_tickets/</code> on the active management console.
OBSOLETE	The ticket creation operation failed during data collection.

## Updating the ticket database when nodes are added or removed

After adding or removing a Management Server or file serving node, run the `/opt/diagnostics/tools/mxdstool addnodes` command on the active management console. This command registers nodes in the ticket database.

## Configuring the support ticket feature

The support ticket feature is typically configured after the X9000 Software installation. To reconfigure this feature, complete the following steps:

1. Configure password-less SSH on X9000 Management Servers (active/passive) and all file serving nodes in the cluster, as described in the following section.
2. Verify that the `/etc/hosts` file on each node contains the hostname entries of all the nodes in the cluster. If not, add them.
3. Run the `/opt/diagnostics/tools/mxdstool addnodes` command manually on the active management console.



**NOTE:**

During the X9000 Software installation, the names of crash dumps in the `/var/crash` directory change to include `_PROCESSED`. For example, `2010-03-08-10:09` changes to `2010-03-08-10:09_PROCESSED`.

**NOTE:**

Be sure to monitor the `/var/crash` directory and remove any unneeded processed crash dumps.

## Configuring shared ssh keys

To configure one-way shared ssh keys on the cluster, complete the following steps:

1. On the management console, run the following commands as root:

```
# mkdir -p $HOME/.ssh
# chmod 0700 $HOME/.ssh
# ssh-keygen -t dsa -f $HOME/.ssh/id_dsa -P ''
```

This command creates two files: `$HOME/.ssh/id_dsa` (private key) and `$HOME/.ssh/id_dsa.pub` (public key).

2. On the management console, run the following command for each file serving node:

```
# ssh-copy-id -i $HOME/.ssh/id_dsa.pub server
```

3. On the Management Console, test the results by using the `ssh` command to connect to each file serving node:

```
# ssh {hostname for file serving node}
```

## Viewing software version numbers

To view version information for a list of hosts, use the following command:

```
<installdirectory>/bin/ibrix_version -l [-h HOSTLIST]
```

For each host, the output includes:

- Version number of the installed file system
- Version numbers of the IAD and File System module
- Operating system type and OS kernel version
- Processor architecture

The `-S` option shows this information for all file serving nodes. The `-C` option shows the information for all X9000 clients.

The file system and IAD/FS output fields should show matching version numbers unless you have installed special releases or patches. If the output fields show mismatched version numbers and you do not know of any reason for the mismatch, contact HP Support. A mismatch might affect the operation of your cluster.

# Troubleshooting specific issues

## Software services

### Cannot start services on the management console, a file serving node, or a Linux X9000 client

SELinux might be enabled. To determine the current state of SELinux, use the `getenforce` command. If it returns `enforcing`, disable SELinux using either of these commands:

```
setenforce Permissive
setenforce 0
```

To permanently disable SELinux, edit its configuration file (`/etc/selinux/config`) and set `SELINUX=parameter` to either `permissive` or `disabled`. SELinux will be stopped at the next boot.

For X9000 clients, the client might not be registered with the management console. For information on registering clients, see the *HP StorageWorks X9000 File Serving Software Installation Guide*.

## Failover

### Cannot fail back from failover caused by storage subsystem failure

When a storage subsystem fails and automated failover is turned on, the management console will initiate its failover protocol. It updates the configuration database to record that segment ownership has transferred from primary servers to their standbys and then attempts to migrate the segments to the standbys. However, segments cannot migrate because neither the primary servers nor the standbys can access the storage subsystem and the failover is stopped.

Perform the following manual recovery procedure:

1. Restore the failed storage subsystem (for example, replace failed Fibre Channel switches or replace a LUN that was removed from the storage array).
2. Reboot the standby servers, which will allow the failover to complete.

### Cannot fail back because of a storage subsystem failure

This issue is similar to the previous issue. If a storage subsystem fails after you have initiated a failback, the configuration database will record that the failback occurred, even though segments never migrated back to the primary server. If you execute `ibrix_fs -i -f FSNAME`, the output will list `No` in the `ONBACKUP` field, indicating that the primary server now owns the segments, even though it does not. In this situation, you will be unable to complete the failback after you fix the storage subsystem problem.

Perform the following manual recovery procedure:

1. Restore the failed storage subsystem.
2. Reboot the primary server, which will allow the arrested failback to complete.

### X9000 client I/O errors following segment migration

Following successful segment migration to a different file serving node, the management console sends all X9000 clients an updated map reflecting the changes, which enables the clients to continue I/O

operations. If, however, the network connection between a client and the management console is not active, the client cannot receive the updated map, resulting in client I/O errors.

To fix the problem, restore the network connection between the clients and the management console.

## Windows X9000 clients

### Logged in but getting a “Permission Denied” message

The X9000 client cannot access the Active Directory server because the domain name was not specified. Reconfigure the Active Directory settings, specifying the domain name (see the *HP StorageWorks X9000 File Serving Software Installation Guide* for more information.).

### Verify button in the Active Directory Settings tab does not work

This issue has the same cause as the above issue.

### Mounted drive does not appear in Windows Explorer

To make a drive appear in Explorer, after mounting it, log off and then log back on, or reboot the machine. You can also open a DOS command window and access the drive manually.

### Mounted drive not visible when using Terminal Server

Refresh the browser's view of the system by logging off and then logging back on.

### X9000 client auto-startup interferes with debugging

The X9000 client is set to start automatically, which can interfere with debugging a Windows X9000 client problem. To prevent this, reboot the machine in safe mode and change the Windows X9000 client service mode to manual, which enables you to reboot without starting the client.

1. Open the Services control manager (**Control Panel > Administrative Tools > Services**).
2. Right-click **IBRIX Client Services** and select **Properties**.
3. Change the startup type to **Manual**, and then click **OK**.
4. Debug the client problem. When finished, switch the Windows X9000 client service back to automatic startup at boot time by repeating these steps and changing the startup type to **Automatic**.

## Synchronizing information on file serving nodes and the configuration database

To maintain access to a file system, file serving nodes must have current information about the file system. HP recommends that you execute `ibrix_health` on a regular basis to monitor the health of this information. If the information becomes outdated on a file serving node, execute `ibrix_dbck -o` to resynchronize the server's information with the configuration database. For information on `ibrix_health`, see “[Monitoring cluster health](#)” on page 52.

---

#### NOTE:

The `ibrix_dbck` command should be used only under the direction of HP Support.

---

To run a health check on a file serving node, use the following command:

```
<installdirectory>/bin/ibrx_health -i -h HOSTLIST
```

If the last line of the output reports `Passed`, the file system information on the file serving node and management console is consistent.

To repair file serving node information, use the following command:

```
<installdirectory>/bin/ibrx_dbck -o -f FSNAME [-h HOSTLIST]
```

To repair information on all file serving nodes, omit the `-h HOSTLIST` argument.

---

# 15 Replacing components

## Customer replaceable components

---

### ⚠ WARNING!

Before performing any of the procedures in this chapter, read the important warnings, precautions, and safety information in [Appendix C](#) and [Appendix D](#).

---

---

### ❗ IMPORTANT:

To avoid unintended consequences, HP recommends that you perform the procedures in this chapter during scheduled maintenance times.

---

---

### ⚠ CAUTION:

Be sure the replacement is available before removing a component or a blanking panel. Open slots dramatically impact airflow and cooling within the device.

For component and cabling diagrams, see [Appendix A](#).

For information about available spare parts, see [Appendix B](#)

---

## Hot-pluggable and non-hot-pluggable components

Before removing any serviceable part, determine whether the part is hot-pluggable or non-hot-pluggable.

- If the component is hot-pluggable, a power shutdown of the device is not required for replacement of the part.
- If the component is not hot-pluggable, the device must be powered down.

## Returning the defective component

In the materials shipped with a CSR component, HP specifies whether the defective component must be returned to HP. In cases where it is required, you must ship the defective part back to HP within a defined period of time, normally five business days. The defective part must be returned with the associated documentation in the provided shipping material. Failure to return the defective part could result in HP billing you for the replacement. With a CSR, HP will pay all shipping and part return costs and determine the courier/carrier to be used.

## Parts-only warranty service

Your HP Limited Warranty could include a parts-only warranty service. Under the terms of parts-only warranty service, HP provides replacement parts, free of charge.

For parts-only warranty service, CSR part replacement is mandatory. If you request HP to replace these parts for you, you are charged for the travel and labor costs of this service.

## Required tools

The following tools might be necessary for some procedures:

- T-10 Torx screwdriver
- T-15 Torx screwdriver
- 4-mm flat-blade screwdriver
- Phillips screwdriver

## Additional documentation

In addition to this document, you will need the following documents, which are available at <http://www.hp.com/support/manuals>:

- *HP ProLiant DL380 G6 Server Maintenance and Service Guide* (for file serving node procedures)
- *HP ProLiant DL360 G6 Server Maintenance and Service Guide* (for X9300 Management Server procedures)

Repair procedures for the X9300 controller and drive enclosure are available in the **Service and maintenance information** section on the HP StorageWorks 2000fc G2 Modular Smart Array Manuals page. To locate the documents, go to <http://www.hp.com/support/manuals>. In the storage section, click **Disk Storage Systems**, and then under MSA Disk Arrays, click **HP StorageWorks 2000fc G2 Modular Smart Array**.

## Replacing a system board

To replace a system board on a file serving node or the X9300 Management Server:

1. If the server is an X9300 Management Server, skip this step. If the server is a file serving node, fail over the server to its standby server using the management console GUI or CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Next, select the server name in the left pane, and click **Failover**.
  - On the CLI, execute the following command, where *server\_name* is the server containing the board:

```
ibrix_server -f -h <server_name>
```
2. Power down the server.
3. Replace the system board: For a file serving node, see the *HP ProLiant DL380 G6 Server Maintenance and Service Guide*. For an X9300 Management Server, see the *HP ProLiant DL360 G6 Server Maintenance and Service Guide*.

4. Power on the server.
  - a. Check the MAC address of the new NIC. Use the `ifdown eth(n)` command, and then use the `ip addr` command.
  - b. Change the `HWADDR=` line in the relevant `ifcfg-eth(n)` file to be the MAC address from the previous step.
  - c. Follow this procedure for every NIC that was affected by the physical board swapout (for example, if this was a dual NIC, then two files need to be updated).
  - d. Restart the network subsystem with `service network restart`.
  - e. Run the following command to verify that all is correct  

```
cat /proc/net/bonding/bond(n)
```
5. For a file serving node, fail back the server using the GUI or CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Next, select the server name in the left pane, and click **Failback**.
  - On the CLI, execute the following command:  

```
ibrix_server -f -U -h <server_name>
```

## Replacing a NIC adapter

To replace a NIC adapter on a file serving node or the X9300 Management Server:

1. If the server is an X9300 Management Server, skip this step. If the server is a file serving node, fail over the server to its standby server using the management console GUI or CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Next, select the server name in the lower-left pane, and click **Failover**.
  - On the CLI, execute the following command, where `server_name` is the server containing the NIC adapter:  

```
ibrix_server -f -h <server_name>
```
2. Power down the server.
3. Replace the NIC adapter.
  - For a file serving node, see the *HP Proliant DL380 G6 Server Maintenance and Service Guide*.
  - For an X9300 Management Server, see the *HP Proliant DL360 G6 Server Maintenance and Service Guide*.
4. Power on the server:
  - a. Check the MAC address of the new NIC. Use the `ifdown eth(n)` command, then use the `ip addr` command.
  - b. Change the `HWADDR=` line in the relevant `ifcfg-eth(n)` file to the MAC address from the previous step.
  - c. Repeat steps 4a and 4b for every NIC that was affected by the physical board swap out (for example, if this was a dual NIC, then two files need to be updated).
  - d. Restart the network subsystem with `service network restart`.
  - e. Verify that all is correct with `cat /proc/net/bonding/bond(n)`.

5. For a file serving node, fail back the server using the GUI or CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Next, select the server name in the left pane, and click **Failback**.
  - On the CLI, execute the following command:

```
ibrix_server -f -U -h <server_name>
```

## Replacing a Fibre Channel HBA

To replace a Fibre Channel HBA on a file serving node:

1. Fail over the server to its standby server using the management console GUI or CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Next, select the server name in the lower-left pane, and click **Failover**.
  - On the CLI, execute the following command, where `server_name` is the server containing the HBA:

```
ibrix_server -f -h <server_name>
```
2. Power down the server.
3. Replace the Fibre Channel HBA.
  - For a file serving node, see the *HP ProLiant DL380 G6 Server Maintenance and Service Guide*.
  - If you are using WWPN/WWNN zoning, write down the WWPN/WWNN of the new HBA and update the zoning.
4. Power on the server.
5. Modify the HBA configuration. If the HBA belonged to a standby pair, delete the pair. If the HBA was monitored, disable monitoring, using the GUI or the CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Select **HBAs** in the left pane, select the appropriate HBA, and then click **Modify**. On the Modify HBA Properties window, deselect the **Enable monitoring and Enable backup** pair.
  - On the CLI, execute the following command to delete a standby pairing:

```
ibrix_hba -b -U -P WWPN1:WWPN2 -h <server_name>
```

Execute the following command on both ports to disable monitoring:

```
ibrix_hba -m -U -h <server_name> -p WWPN
```
6. Delete the HBA from the cluster configuration database using the GUI or CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Select **HBAs** in the left pane. select an HBA, and click **Delete**.
  - On the CLI, execute the following command for both HBA ports:

```
ibrix_hba -d -h <server_name> -w WWNN
```
7. Discover the new HBA ports using the GUI or CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Select **HBAs** in the left pane and click **Discover**.
  - On the CLI, execute the following command:

```
ibrix_hba -a -h <server_name>
```



8. Enable standby pairing and monitoring on the new HBA if desired, using the GUI or CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Select **HBAs** in the left pane, select the appropriate Node WWN, and click **Modify**. On the Modify HBA Properties window, select **Enable monitoring and Enable backup** pair, and select the WWPN for the backup HBA.
  - On the CLI, execute the following command to enable standby pairing:  

```
ibrix_hba -b -P WWPN1:WWPN2 -h <server_name>
```

Execute the following command on both ports to enable monitoring:

```
ibrix_hba -m -h <server_name> -p WWPN
```
9. Fail back the server using the GUI or CLI:
  - On the GUI, select **Servers** from the Navigator pane, and then select the appropriate server from the Servers pane. Next, select the server name in the left pane, and click **Failback**.
  - On the CLI, execute the following command:  

```
ibrix_server -f -U -h <server_name>
```



---

# 16 Recovering a file serving node

Use the following procedure to recover a failed file serving node. You will need to create a QuickRestore DVD, as described later, and then install it on the affected node. This step installs the operating system and X9000 Software on the node and launches a configuration wizard. Note the following:

- If you are restoring a file serving node that was previously a host for the agile management console, you can configure the node as a passive agile management console during the recovery.
- If you are restoring a Management Server that was migrated to be a host for the agile management console, start the recovery, as described in the following section, and then use the manual procedure for configuring a file serving node. Configure the Management Server as a passive agile management console.

---

△ **CAUTION:**

The Quick Restore DVD restores the server to its original factory state. This is a destructive process that completely erases all of the data on local hard drives.

---

## Starting the recovery

To recover a failed file serving node, follow these steps:

1. If a NIC monitor is configured on the user network, remove the monitor. To determine if NIC monitoring is configured, run the following command on the Management Server:

```
ibrix_nic -i -h <hostname>
```

Check the output for a line such as the following:

```
Monitored By : titan16
```

To remove the monitor, use the following command:

```
ibrix_nic -m -h MONITORHOST -D DESTHOST/IFNAME
```

For example:

```
ibrix_nic -m -h titan16 -D titan15/eth2
```

2. Obtain the latest Quick Restore image from the HP kiosk <http://www.software.hp.com/kiosk> (you will need your HP-provided login credentials).
3. Burn the ISO image to a DVD.
4. Insert the Quick Restore DVD into the server's DVD-ROM drive.
5. Restart the server to boot from the DVD-ROM.

6. When the following screen appears, enter **qr** to recover the file serving node.



The server reboots automatically after the software is installed. Remove the DVD from the DVD-ROM drive.

7. When your cluster was configured initially, the installer may have created a template for configuring file serving nodes. To use this template to configure the file serving node undergoing recovery, go to [“Configuring a file serving node using the original template”](#) on page 108.

To configure the file serving node manually, without the template, go to [“Configuring a file serving node manually”](#) on page 112.

To configure a Management Server that hosts the agile management console, go to [“Configuring a file serving node manually”](#) on page 112.


## Configuring a file serving node using the original template

Complete the following steps:

1. Log into the system as user `root` (the default password is `hpinvent`).
2. When the System Deployment Menu appears, select **Join an existing cluster**.



3. The Configuration Wizard attempts to discover management consoles on the network and then displays the results. Select the appropriate management console for this cluster.



**NOTE:**

If the list does not include the appropriate management console, or you want to customize the cluster configuration for the file serving node, select **Cancel**. Go to [Configuring a file serving node manually](#), page 112 for information about completing the configuration.

4. On the Verify Hostname dialog box, enter a hostname for the node, or accept the hostname generated by the hostname template.



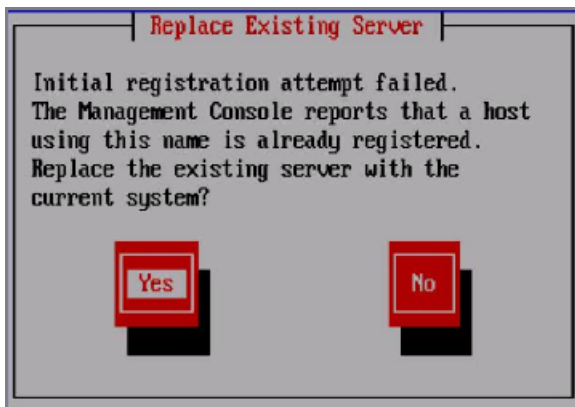
5. The Verify Configuration window shows the configuration received from the management console. Select **Accept** to apply the configuration to the server and register the server with the management console.



#### NOTE:

If you select **Reject**, the wizard will exit and the shell prompt will be displayed. You can restart the Wizard by entering the command `/usr/local/ibrix/autocfg/bin/menu_ss_wizard` or logging in to the server again.

6. If the specified hostname already exists in the cluster (the name was used by the node you are replacing), the Replace Existing Server window asks whether you want to replace the existing server with the node you are configuring. When you click **Yes**, the replacement node will be registered.

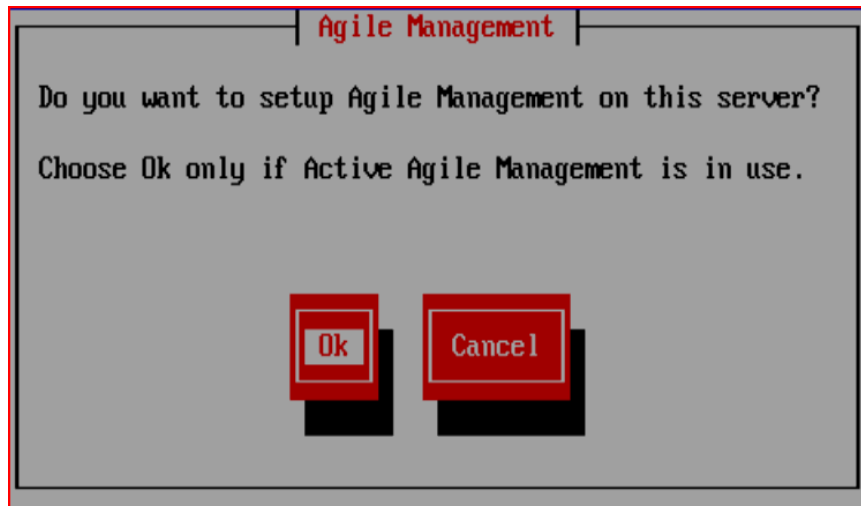


7. After the server is registered, you will be asked whether you want to configure a passive agile management console on the node.

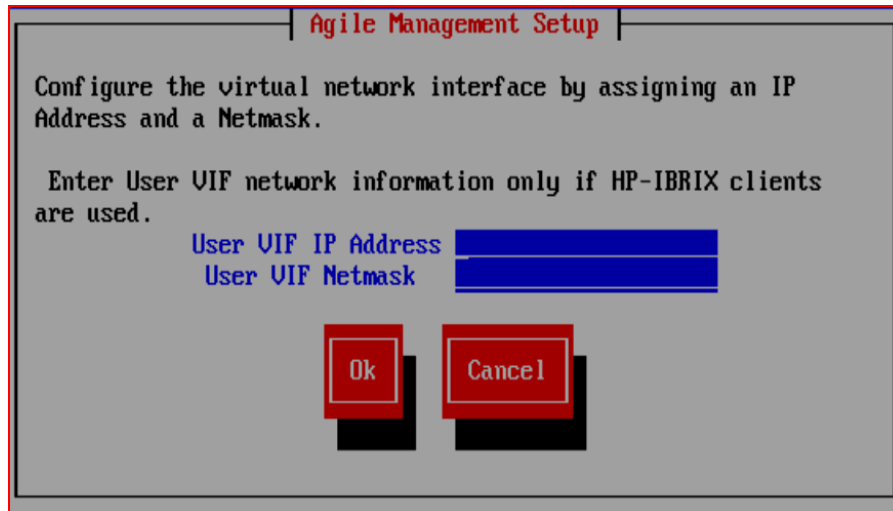
❗ **IMPORTANT:**

Configure a passive agile management console only if the following conditions are met:

- The agile management console is enabled and an active agile management console is configured.
- The passive agile management console is not configured on any other nodes in the cluster.



If you have configured a user network, enter a VIF IP address and netmask.



If you configured a passive management console, enter the following command to verify the status of the console:

**ibrix\_fm -i**

Next, complete the restore on the file serving node.

# Completing the restore on a file serving node

Complete the following steps:

1. Ensure that you have root access to the node. The restore process sets the root password to **hpinvent**, the factory default.
2. The QuickRestore DVD enables the `iptables` firewall. Either make the firewall configuration match that of your other server blades to allow traffic on appropriate ports, or disable the service entirely by running the `chkconfig iptables off` and `service iptables stop` commands.

To allow traffic on appropriate ports, open the following ports:

- 80
- 443
- 1234
- 9000
- 9005
- 9008
- 9009

3. If you disabled NIC monitoring before using the QuickRestore DVD, re-enable the monitor:

```
ibrix_nic -m -h MONITORHOST -A DESTHOST/IFNAME
```

For example:

```
ibrix_nic -m -h titan16 -A titan15/eth2
```

4. Run `ibrix_health -l` from the X9000 management console to verify that no errors are being reported.



## NOTE:

When you perform a Quick Restore of a file serving node, the NFS, CIFS, FTP, and HTTP export information is not automatically restored to the node. After operations are failed back to the node, the I/O from client systems to the node fails for the NFS, CIFS, FTP, and HTTP shares. To avoid this situation, manually restore the NFS, CIFS, FTP, and HTTP exports on the node before failing it back.

# Configuring a file serving node manually

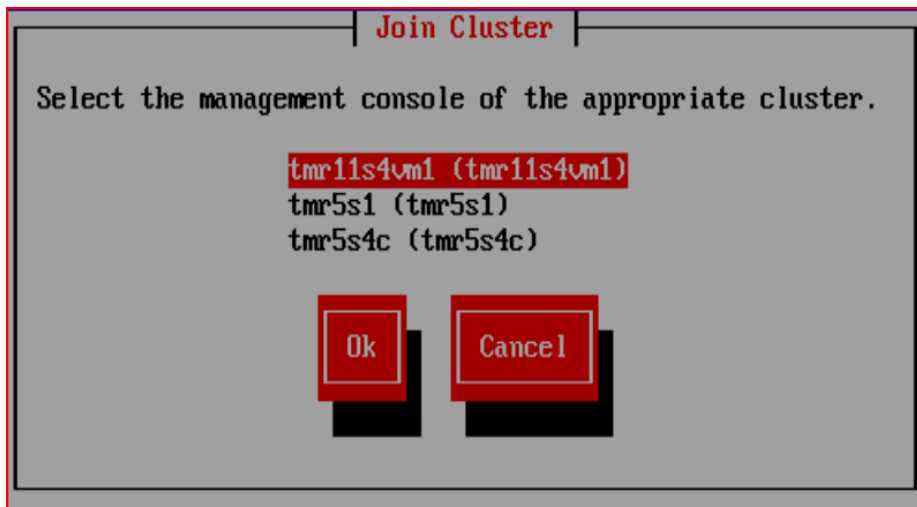
Use this procedure to configure file serving nodes manually instead of using the template. Also use this procedure to configure a Management Server that acts as a host for the agile management console. (You can launch the wizard manually by entering the command `/usr/local/ibrix/autocfg/bin/menu_ss_wizard`.)

1. Log into the system as user `root` (the default password is `hpinvent`).
2. When the System Deployment Menu appears, select **Join an existing cluster**.

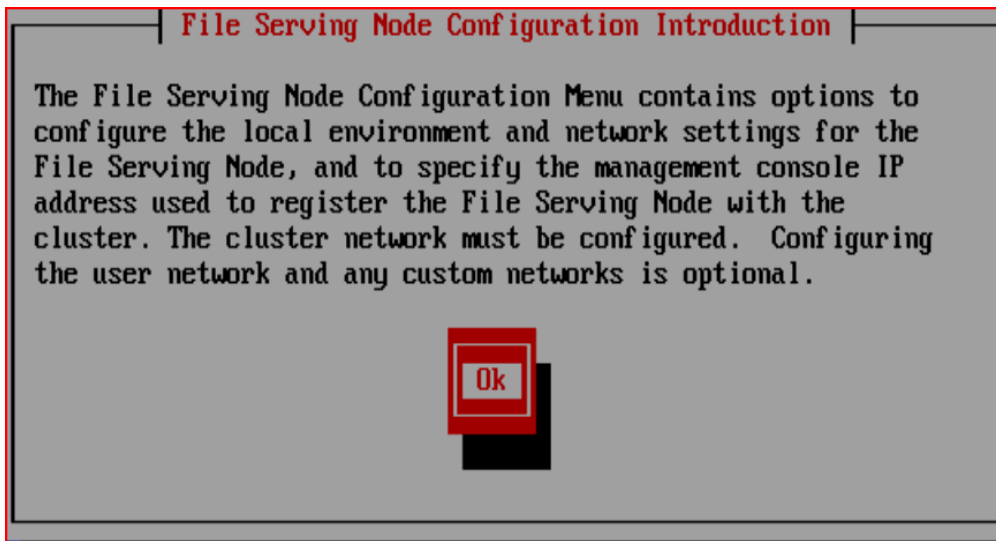




3. The Configuration Wizard attempts to discover management consoles on the network and then displays the results. Select **Cancel** to configure the node manually.



4. The file serving node Configuration Menu appears.




5. The Cluster Configuration Menu lists the configuration parameters that you will need to set. Use the Up and Down arrow keys to select an item in the list. When you have made your select, press **Tab** to move to the buttons at the bottom of the dialog box, and press **Space** to go to the next dialog box.



6. Select **Management Console** from the menu, and enter the IP address of the management console. This is typically the address of the management console on the cluster network.




7. Select **Hostname** from the menu, and enter the hostname of this server.



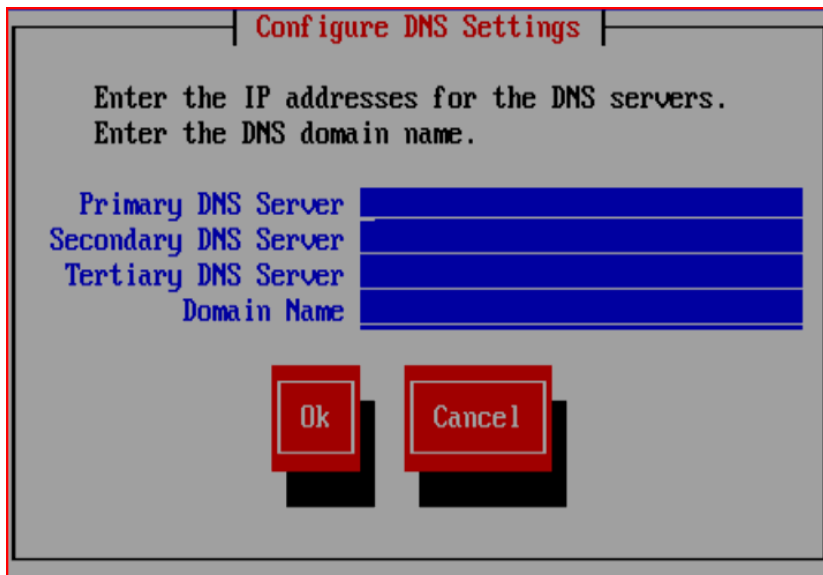
8. Select **Time Zone** from the menu, and then use **Up** or **Down** to select your time zone.



9. Select **Default Gateway** from the menu, and enter the IP Address of the host that will be used as the default gateway.

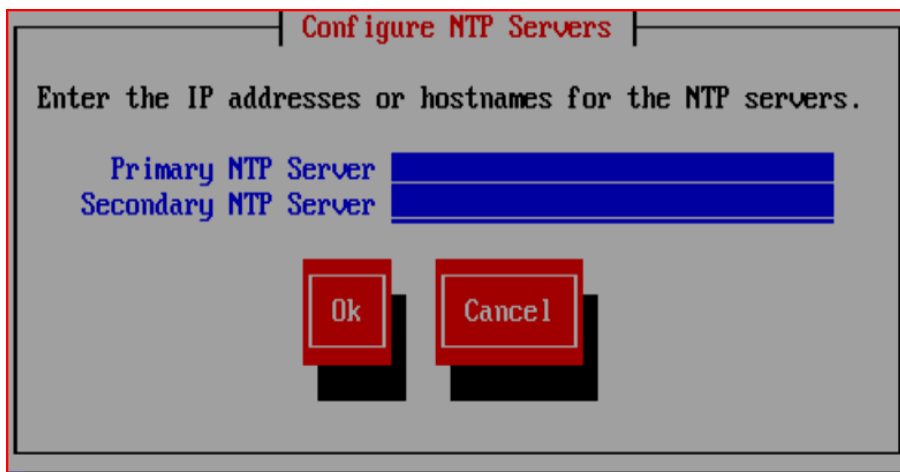


10. Select **DNS Settings** from the menu, and enter the IP addresses for the primary and secondary DNS servers that will be used to resolve domain names. Also enter the DNS domain name.



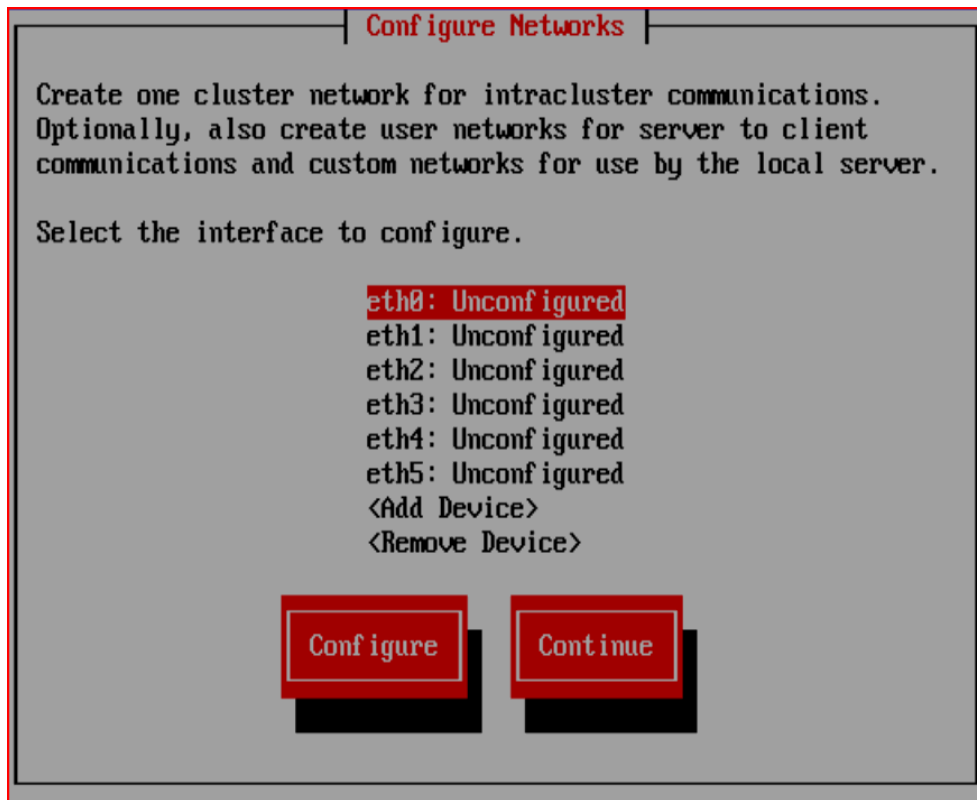
A screenshot of a 'Configure DNS Settings' dialog box. The title bar is red with the text 'Configure DNS Settings' in white. The main area has a gray background. It contains the text 'Enter the IP addresses for the DNS servers.' and 'Enter the DNS domain name.' in black. Below this, there are four blue input fields with labels: 'Primary DNS Server', 'Secondary DNS Server', 'Tertiary DNS Server', and 'Domain Name'. At the bottom, there are two red buttons with white text: 'Ok' and 'Cancel'.

11. Select **NTP Servers** from the menu, and enter the IP addresses or hostnames for the primary and secondary NTP servers.




A screenshot of a 'Configure NTP Servers' dialog box. The title bar is red with the text 'Configure NTP Servers' in white. The main area has a gray background. It contains the text 'Enter the IP addresses or hostnames for the NTP servers.' in black. Below this, there are two blue input fields with labels: 'Primary NTP Server' and 'Secondary NTP Server'. At the bottom, there are two red buttons with white text: 'Ok' and 'Cancel'.

12. Select **Networks** from the menu. Select **<add device>** to create a bond for the cluster network.



You are creating a bonded interface for the cluster network; select **Ok** on the Select Interface Type dialog box.



Enter a name for the interface (bond0 for the cluster interface) and specify the appropriate options and slave devices.

Add Bonded Interface

Enter a name for the interface such as bond0 or bond1. Specify the appropriate bond options and the select the slave devices.

Bond Name

Bond Options

Slave Devices

mode=6 miimon=100 updelay=100\_\_

eth1

eth2

eth3

eth4

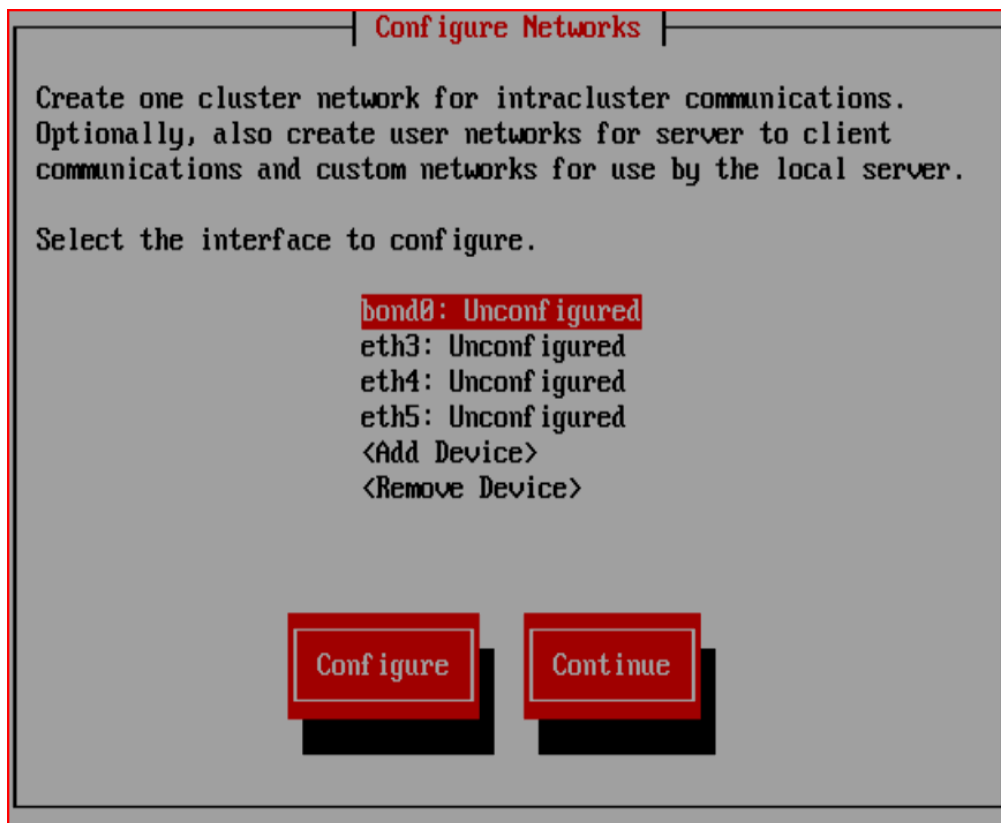
eth5

Ok

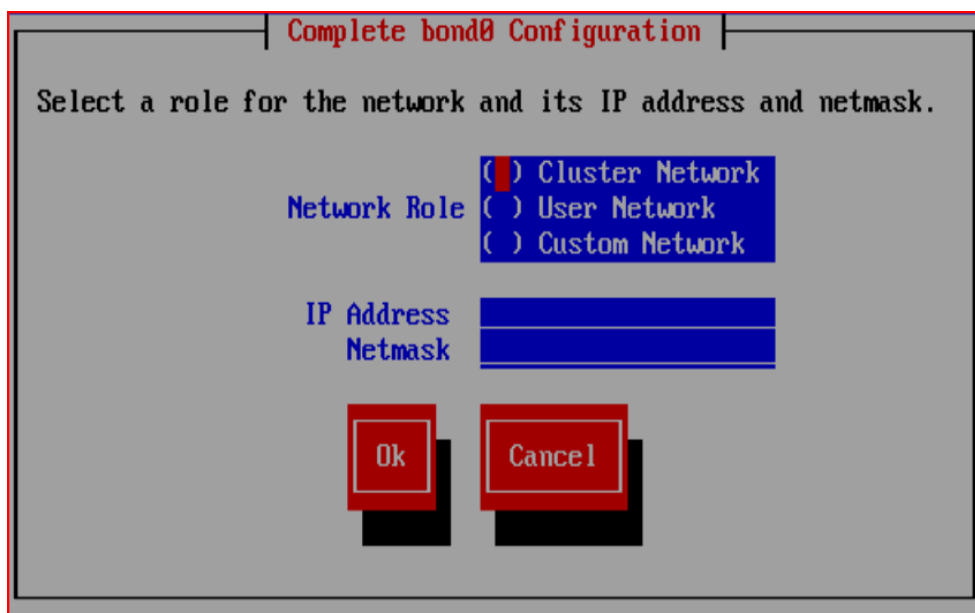
Cancel

13. When the Configure Network dialog box reappears, select bond0.

X9300 Network Storage Gateway Administrator Guide 119



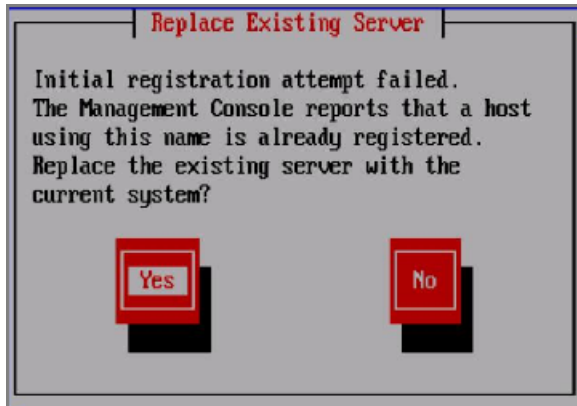
14. To complete the `bond0` configuration, enter a space to select the Cluster Network role. Then enter the IP address and netmask information that the network will use.



Repeat this procedure to create a bonded user network (typically `bond1`) and any custom networks as required.



15. When you have completed your entries on the file serving node Configuration Menu, select **Continue**.
16. Verify your entries on the confirmation screen, and select **Commit** to apply the values to the file serving node and register it with the management console.
17. If the hostname specified for the node already exists in the cluster (the name was used by the node you are replacing), the Replace Existing Server window asks whether you want to replace the existing server with the node you are configuring. When you click **Yes**, the replacement node will be registered.



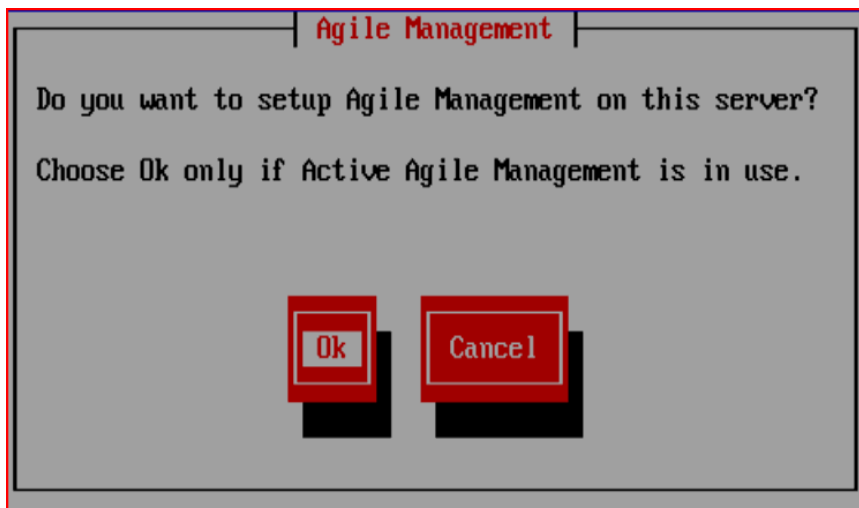
18. After the server is registered, you will be asked whether you want to configure a passive agile management console on the node.

---

❗ **IMPORTANT:**

Configure a passive agile management console only if the following conditions are met:

- The agile management console is enabled and an active agile management console is configured.
  - The passive agile management console is not configured on any other nodes in the cluster.
- 



If you configured a user network, enter a VIF IP address and netmask for the network.

Agile Management Setup

Configure the virtual network interface by assigning an IP Address and a Netmask.

Enter User VIF network information only if HP-IBRIX clients are used.

User VIF IP Address

User VIF Netmask

Ok

Cancel

If you configured a passive management console, enter the following command to verify the status of the console:

**ibrix\_fm -i**

Next, go to [Completing the restore on a file serving node](#), page 112.

---

# 17 Support and other resources

## Contacting HP

For worldwide technical support information, see the HP support website:

<http://www.hp.com/support>

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

## Related information

The following documents [and websites] provide related information:

- HP StorageWorks X9300 Network Storage Gateway documentation:
  - *HP StorageWorks X9000 File Serving Software User Guide*
  - *HP StorageWorks X9000 File Serving Software CLI Reference*
  - *HP StorageWorks X9000 File Serving Software Release Notes*

To locate the X9300 Network Storage Gateway documents, go to <http://www.hp.com/go/X9000>.

- HP StorageWorks X9300 Management Server documentation:
  - *HP ProLiant DL360 G6 Server User Guide*
  - *HP ProLiant DL360 G6 Server Maintenance and Service Guide*

To locate the HP StorageWorks X9300 Management Server documents, go to <http://www.hp.com/support/manuals>,

In the servers section, click **ProLiant ml/dl and tc series servers**, and then click **HP ProLiant DL360 G6 Server series**.

- HP StorageWorks X9300 file serving node documentation:
  - *HP ProLiant DL380 G6 Server User Guide*
  - *HP ProLiant DL380 G6 Server Maintenance and Service Guide*

To locate the HP StorageWorks X9300 file serving node documents, go to <http://www.hp.com/support/manuals>,

In the servers section, click **ProLiant ml/dl and tc series servers**, and then click **HP ProLiant DL380 G6 Server series**.

## HP websites

For additional information, see the following HP websites:

- <http://www.hp.com/go/X9000>
- <http://www.hp.com>
- <http://www.hp.com/go/storage>
- [http://www.hp.com/service\\_locator](http://www.hp.com/service_locator)
- <http://www.hp.com/support/manuals>
- <http://www.hp.com/support/downloads>
- <http://www.hp.com/storage/whitepapers>

## Rack stability

Rack stability protects personnel and equipment.

---

### WARNING!

To reduce the risk of personal injury or damage to equipment:

- Extend leveling jacks to the floor.
  - Ensure that the full weight of the rack rests on the leveling jacks.
  - Install stabilizing feet on the rack.
  - In multiple-rack installations, fasten racks together securely.
  - Extend only one rack component at a time. Racks can become unstable if more than one component is extended.
- 

## Customer self repair

HP customer self repair (CSR) programs allow you to repair your StorageWorks product. If a CSR part needs replacing, HP ships the part directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your HP-authorized service provider will determine whether a repair can be accomplished by CSR.


For more information about CSR, contact your local service provider, or see the CSR website:

<http://www.hp.com/go/selfrepair>

# A Component and cabling diagrams


## Component diagrams

### Front view of file serving node



Item	Description
1	Quick-release levers (2)
2	HP Systems Insight Manager display
3	Hard drive bays
4	SATA optical drive bay
5	Video connector
6	USB connectors (2)

### Rear view of file serving node

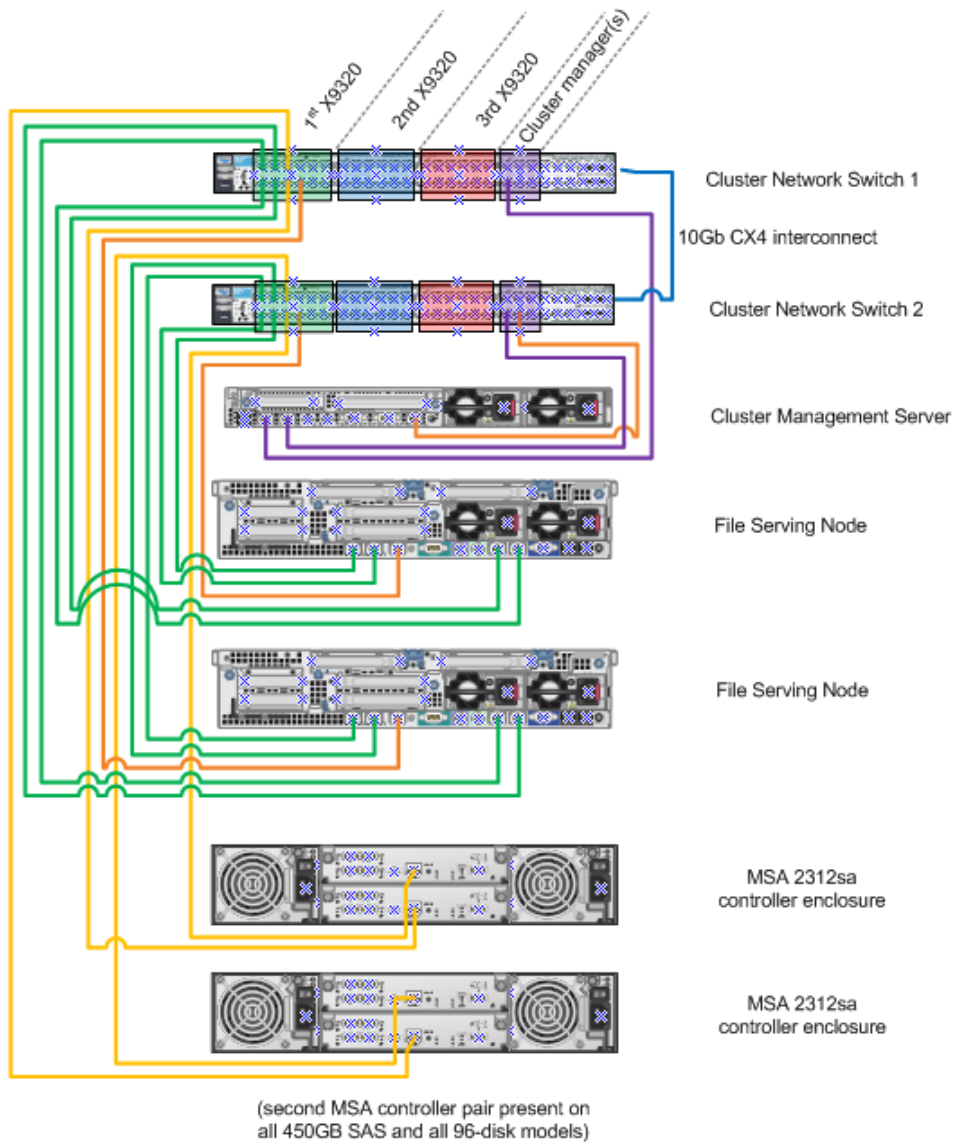


Item	Description
1	PCI slot 5
2	PCI slot 6
3	PCI slot 4
4	PCI slot 2
5	PCI slot 3
6	PCI slot 1
7	Power supply 2 (PS2)
8	Power supply 1 (PS1)
9	USB connectors (2)
10	Video connector
11	NIC 1 connector
12	NIC 2 connector
13	Mouse connector
14	Keyboard connector
15	Serial connector
16	iLO 2 connector
17	NIC 3 connector
18	NIC 4 connector

Server	PCIe card	PCI slot
SATA 1Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	NC364T Quad 1Gb NIC	2
	empty	3
	empty	4
	empty	5
	empty	6
SATA 10Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	empty	2
	empty	3
	NC522SFP dual 10Gb NIC	4
	empty	5
	empty	6
SAS 1Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	NC364T Quad 1Gb NIC	2
	empty	3
	HP SC08Ge 3Gb SAS Host Bus Adapter	4
	empty	5
	empty	6
SAS 10Gb	HP SC08Ge 3Gb SAS Host Bus Adapter	1
	HP SC08Ge 3Gb SAS Host Bus Adapter	2
	empty	3
	NC522SFP dual 10Gb NIC	4
	empty	5
	empty	6

# Cabling diagrams

## Cluster network cabling diagram



### NOTE:

The MSA controllers are optional in the X9300 models.



# B Spare parts list

This appendix lists spare parts (both customer replaceable and non customer replaceable) for the X9300 Network Storage Gateway components.

Spare parts are categorized as follows:

- **Mandatory.** Parts for which customer self repair is mandatory. If you request HP to replace these parts, you will be charged for the travel and labor costs of this service.
- **Optional.** Parts for which customer self repair is optional. These parts are also designed for customer self-repair. If, however, you require that HP replace them for you, there may or may not be additional charges, depending on the type of warranty service designated for your product.



## NOTE:

Some HP parts are not designed for customer self-repair. In order to satisfy the customer warranty, HP requires that an authorized service provider replace the part. These parts are identified as "No" in the spare parts lists.

## 1GbE spare parts

### 1 GbE (AW539A)

Description	Spare part number	Customer self repair
SPS-HOOD	496056-001	Mandatory
SPS-CARD, RISER	496057-001	Optional
SPS-KIT, MISC HARDWARE	496058-001	Mandatory
SPS-CAGE, PCI	496060-001	Optional
SPS-BAFFLE	496061-001	Mandatory
SPS-BACKPLANE, PS	496062-001	Optional
SPS-CAGE, PS BACKPLANE	496063-001	Optional
SPS-HEATSINK, 80W	496064-001	Optional
SPS-FAN	496066-001	Mandatory
SPS-CAGE, FAN	496067-001	Mandatory
SPS-BD,SYSTEM I/O, W/SUBPAN	496069-001	Optional

Description	Spare part number	Customer self repair
SPS-CABLE, SAS BACKPLANE	496070-001	Mandatory
SPS-CABLE, SATA DVD PWR	496071-001	Mandatory
SPS-BD, SID	496073-001	Mandatory
SPS-CAGE, HD, SFF	496074-001	Mandatory
SPS-CAGE, DVD OPT DRIVE	496076-001	Mandatory
SPS-BD, PCIX	496077-001	Optional
SPS-BD, PCIE	496078-001	Optional
SPS-BEZEL	496080-001	Mandatory
SPS-BACKPLANE,SAS	507690-001	Optional
SPS-PROC,NEHALEM EP 2.26 GHZ, 8M, 80W	490073-001	Optional
SPS-DIMM,8GB PC3-10600R,512MX4,ROHS	501536-001	Optional
SPS-DRV,HD,146GB,15K 2.5" DP HP SAS	504334-001	Mandatory
SPS - HW PLASTICS KIT DL180 G6	507260-001	Mandatory
SPS-DRV,ODD, SLIM SATA DVD RW	481429-001	Optional
SPS-TRAY, DVD	532390-001	Mandatory
SPS-POWER SUPPLY, 750W	511778-001	Optional
SPS-BD,NIC,X4 PCI-E,4 PORT,1000 BASE-T	436431-001	Mandatory
SPS-HARDWARE MTG KIT	574765-001	Mandatory
SPS-CORD,AC PWR IEC/IEC 6 FT	142258-001	Mandatory

## 1 GbE (AW539B)

Description	Spare part number	Customer self repair
SPS-HOOD	496056-001	Mandatory
SPS-CARD, RISER	496057-001	Optional
SPS-KIT, MISC HARDWARE	496058-001	Mandatory
SPS-CAGE, PCI	496060-001	Optional
SPS-BAFFLE	496061-001	Mandatory
SPS-BACKPLANE, PS	496062-001	Optional

<b>Description</b>	<b>Spare part number</b>	<b>Customer self repair</b>
SPS-CAGE, PS BACKPLANE	496063-001	Optional
SPS-HEATSINK, 80W	496064-001	Optional
SPS-FAN	496066-001	Mandatory
SPS-CAGE, FAN	496067-001	Mandatory
SPS-BD,SYSTEM I/O, W/SUBPAN	496069-001	Optional
SPS-CABLE, SAS BACKPLANE	496070-001	Mandatory
SPS-CABLE, SATA DVD PWR	496071-001	Mandatory
SPS-BD, SID	496073-001	Mandatory
SPS-CAGE, HD, SFF	496074-001	Mandatory
SPS-CAGE, DVD OPT DRIVE	496076-001	Mandatory
SPS-BD, PCIX	496077-001	Optional
SPS-BD, PCIE	496078-001	Optional
SPS-BEZEL	496080-001	Mandatory
SPS-BACKPLANE,SAS	507690-001	Optional
SPS-PROC,NEHALEM EP 2.26 GHZ, 8M, 80W	490073-001	Optional
SPS-DIMM,4GB PC3-10600R,256MX4,ROHS	501534-001	Mandatory
SPS-DRV,HD,146GB,10K 2.5" DP HP 6G SAS	507283-001	Mandatory
SPS - HW PLASTICS KIT DL180 G6	507260-001	Mandatory
SPS-DRV,ODD, SLIM SATA DVD RW	481429-001	Optional
SPS-TRAY, DVD	532390-001	Mandatory
SPS-POWER SUPPLY, 750W	511778-001	Optional
SPS-BD,NIC,X4 PCI-E,4 PORT,1000 BASE-T	436431-001	Mandatory
SPS-HARDWARE MTG KIT	574765-001	Mandatory
SPS-CORD,AC PWR IEC/IEC 6 FT	142258-001	Mandatory

# 10 GbE spare parts

## 10 GbE (AW540A)

Description	Spare part number	Customer self repair
SPS-HOOD	496056-001	Mandatory
SPS-CARD, RISER	496057-001	Optional
SPS-KIT, MISC HARDWARE	496058-001	Mandatory
SPS-CAGE, PCI	496060-001	Optional
SPS-BAFFLE	496061-001	Mandatory
SPS-BACKPLANE, PS	496062-001	Optional
SPS-CAGE, PS BACKPLANE	496063-001	Optional
SPS-HEATSINK, 80W	496064-001	Optional
SPS-FAN	496066-001	Mandatory
SPS-CAGE, FAN	496067-001	Mandatory
SPS-BD,SYSTEM I/O, W/SUBPAN	496069-001	Optional
SPS-CABLE, SAS BACKPLANE	496070-001	Mandatory
SPS-CABLE, SATA DVD PWR	496071-001	Mandatory
SPS-BD, SID	496073-001	Mandatory
SPS-CAGE, HD, SFF	496074-001	Mandatory
SPS-CAGE, DVD OPT DRIVE	496076-001	Mandatory
SPS-BD, PCIX	496077-001	Optional
SPS-BD, PCIE	496078-001	Optional
SPS-BEZEL	496080-001	Mandatory
SPS-BACKPLANE,SAS	507690-001	Optional
SPS-PROC,NEHALEM EP 2.26 GHZ, 8M, 80W	490073-001	Optional
SPS-DIMM,8GB PC3-10600R,512MX4,ROHS	501536-001	Optional
SPS-DRV,HD,146GB,15K 2.5" DP HP SAS	504334-001	Mandatory
SPS - HW PLASTICS KIT DL180 G6	507260-001	Mandatory
SPS-DRV,ODD, SLIM SATA DVD RW	481429-001	Optional

Description	Spare part number	Customer self repair
SPS-TRAY, DVD	532390-001	Mandatory
SPS-POWER SUPPLY, 750W	511778-001	Optional
SPS-BD,NC522SFP+ 10 GIGABIT,SERVER ADPTR	468349-001	Optional
SPS-HARDWARE MTG KIT	574765-001	Mandatory
SPS-CORD,AC PWR IEC/IEC 6 FT	142258-001	Mandatory

## 10 GbE/IB (AW540B)

Description	Spare part number	Customer self repair
SPS-HOOD	496056-001	Mandatory
SPS-CARD, RISER	496057-001	Optional
SPS-KIT, MISC HARDWARE	496058-001	Mandatory
SPS-CAGE, PCI	496060-001	Optional
SPS-BAFFLE	496061-001	Mandatory
SPS-BACKPLANE, PS	496062-001	Optional
SPS-CAGE, PS BACKPLANE	496063-001	Optional
SPS-HEATSINK, 80W	496064-001	Optional
SPS-FAN	496066-001	Mandatory
SPS-CAGE, FAN	496067-001	Mandatory
SPS-BD,SYSTEM I/O, W/SUBPAN	496069-001	Optional
SPS-CABLE, SAS BACKPLANE	496070-001	Mandatory
SPS-CABLE, SATA DVD PWR	496071-001	Mandatory
SPS-BD, SID	496073-001	Mandatory
SPS-CAGE, HD, SFF	496074-001	Mandatory
SPS-CAGE, DVD OPT DRIVE	496076-001	Mandatory
SPS-BD, PCIX	496077-001	Optional
SPS-BD, PCIE	496078-001	Optional
SPS-BEZEL	496080-001	Mandatory
SPS-BACKPLANE,SAS	507690-001	Optional

Description	Spare part number	Customer self repair
SPS-PROC,NEHALEM EP 2.26 GHZ, 8M, 80W	490073-001	Optional
SPS-DIMM,4GB PC3-10600R,256MX4,ROHS	501534-001	Mandatory
SPS-DRV,HD,146GB,10K 2.5" DP HP 6G SAS	507283-001	Mandatory
SPS - HW PLASTICS KIT DL180 G6	507260-001	Mandatory
SPS-DRV,ODD, SLIM SATA DVD RW	481429-001	Optional
SPS-TRAY, DVD	532390-001	Mandatory
SPS-POWER SUPPLY, 750W	511778-001	Optional
SPS-HARDWARE MTG KIT	574765-001	Mandatory
SPS-CORD,AC PWR IEC/IEC 6 FT	142258-001	Mandatory

## IB (AW541A)

Description	Spare part number	Customer self repair
MandatorySPS-HOOD	SPS-CAGE, PCI496056-001	496060-001Mandatory
SPS-CAGE, PCI	496060-001	Optional
SPS-BAFFLE	496061-001	Mandatory
SPS-BACKPLANE, PS	496062-001	Optional
SPS-CAGE, PS BACKPLANE	496063-001	Optional
SPS-HEATSINK, 80W	496064-001	Optional
SPS-FAN	496066-001	Mandatory
SPS-CAGE, FAN	496067-001	Mandatory
SPS-BD,SYSTEM I/O, W/SUBPAN	496069-001	Optional
SPS-CABLE, SAS BACKPLANE	496070-001	Mandatory
SPS-CABLE, SATA DVD PWR	496071-001	Mandatory
SPS-BD, SID	496073-001	Mandatory
SPS-CAGE, HD, SFF	496074-001	Mandatory
SPS-CAGE, DVD OPT DRIVE	496076-001	Mandatory
SPS-BD, PCIX	496077-001	Optional
SPS-BD, PCIE	496078-001	Optional

Description	Spare part number	Customer self repair
SPS-BEZEL	496080-001	Mandatory
SPS-BACKPLANE,SAS	507690-001	Optional
SPS-PROC,NEHALEM EP 2.26 GHZ, 8M, 80W	490073-001	Optional
SPS-DIMM,8GB PC3-10600R,512MX4,ROHS	501536-001	Optional
SPS-DRV,HD,146GB,15K 2.5" DP HP SAS	504334-001	Mandatory
SPS - HW PLASTICS KIT DL180 G6	507260-001	Mandatory
SPS-DRV,ODD, SLIM SATA DVD RW	481429-001	Optional
SPS-TRAY, DVD	532390-001	Mandatory
SPS-POWER SUPPLY, 750W	511778-001	Optional
SPS-KIT, MISC HARDWARE	496058-001	Mandatory
SPS-BD,4X QDR,PCIE,G2,DUAL PORT	519132-001	Optional
SPS-HARDWARE MTG KIT	574765-001	Mandatory
SPS-CORD,AC PWR IEC/IEC 6 FT	142258-001	Mandatory
SPS-CARD, RISER	496057-001	Optional

## Base rack (AW546A)

Description	Spare part number	Customer self repair
SPS-CORD,AC PWR IEC/IEC 6 FT	142258-001	Mandatory
SPS-BRACKETS,PDU	252641-001	Optional
SPS-SWITCH,SVR CNSL,KVM,0X1X8	340386-001	Optional
SPS-HDW,MNT KIT,CNSL,KVM SWT	341519-001	Optional
SPS-RACK,UNIT,10642,10KG2	385969-001	Mandatory
SPS-PANEL,SIDE,10642,10KG2	385971-001	Mandatory
SPS-STABILIZER,600MM,10KG2	385973-001	Mandatory
SPS-SHOCK PALLET,600MM,10KG2	385976-001	Mandatory
SPS-HARDWARE KIT,10KG2	385978-001	Mandatory
SPS-SWITCH,SVR CNSL,KVM,0X1X8	396630-001	Optional
SPS- CA,SRL/DWNLD,9PIN M/F 6'	397641-001	Not customer replaceable

Description	Spare part number	Customer self repair
SPS-SWITCH,SVR CNSL,KVM,0X2X16,USB	410529-001	Mandatory
SPS-RACK,BUS BAR & WIRE TRAY	457015-001	Optional
SPS-STICK,4X FIXED,C-13,OFFSET,WW	483915-001	Optional
SPS-ASSY,RETENTION,PWR CD, LT&RT BRKT	490992-001	Mandatory
SPS-SWITCH,KVM,SVR CNSL,0X2X16	517691-001	Optional
CABLE, CONSOLE D-SUB9 - RJ45 L250	5188-6699	Mandatory
PWR-CORD OPT-918 3-COND 2.0-M-LG ROHS	8120-4753	Not customer replaceable
PWR-CORD OPT-927 3-COND 2.5-M-LG ROHS	8121-0673	Mandatory
PWR-CORD OPT-903 3-COND 3.0-M-LG ROHS	8121-0822	Not customer replaceable
PWR-CORD OPT-902 3-COND 3.0-M-LG ROHS	8121-0823	Not customer replaceable
PWR-CORD OPT-900 3-COND 3.0-M-LG ROHS	8121-0824	Not customer replaceable
PWR-CORD OPT-912 3-COND 3.0-M-LG ROHS	8121-0826	Not customer replaceable
PWR-CORD OPT-906 3-COND 3.0-M-LG ROHS	8121-0827	Mandatory
PWR-CORD OPT-901 3-COND 3.0-M-LG ROHS	8121-0828	Not customer replaceable
PWR-CORD OPT-922 3-COND 3.0-M-LG ROHS	8121-0829	Not customer replaceable
PWR-CORD OPT-917 3-COND 3-M-LG ROHS	8121-0919	Mandatory
PWR-CORD OPT-934 3-COND 3.6-M-LG ROHS	8121-0965	Mandatory
PWR-CORD OPT-919 3-COND 2.5M-LG ROHS	8121-1035	Mandatory
KIT, 2910-24G SMO Support	J9145-61001	Mandatory
HP PROCURVE 2910AL-24G SWITCH	J9145-69001	Mandatory



---

# C Warnings and precautions

## Electrostatic discharge information

See [Electrostatic discharge](#).

## Grounding methods

There are several methods for grounding. Use one or more of the following methods when handling or installing electrostatic sensitive parts:

- Use a wrist strap connected by a ground cord to a grounded workstation or computer chassis. Wrist straps are flexible straps with a minimum of 1 megohm  $\pm 10$  percent resistance in the ground cords. To provide proper ground, wear the strap snug against the skin.
- Use heel straps, toe straps, or boot straps at standing workstations. Wear the straps on both feet when standing on conductive floors or dissipating floor mats.
- Use conductive field service tools.
- Use a portable field service kit with a folding static-dissipating work mat.

If you do not have any of the suggested equipment for proper grounding, have an HP-authorized reseller install the part.



### NOTE:

For more information on static electricity or assistance with product installation, contact your HP-authorized reseller.

---

## Equipment symbols

If the following symbols are located on equipment, hazardous conditions could exist.



### WARNING!




Any enclosed surface or area of the equipment marked with these symbols indicates the presence of electrical shock hazards. Enclosed area contains no operator serviceable parts. To reduce the risk of injury from electrical shock hazards, do not open this enclosure.

---

---

⚠ **WARNING!**



Any RJ-45 receptacle marked with these symbols indicates a network interface connection. To reduce the risk of electrical shock, fire, or damage to the equipment, do not plug telephone or telecommunications connectors into this receptacle.

---

---

⚠ **WARNING!**



Any surface or area of the equipment marked with these symbols indicates the presence of a hot surface or hot component. Contact with this surface could result in injury.

---

---

⚠ **WARNING!**



Power supplies or systems marked with these symbols indicate the presence of multiple sources of power.

---

---

⚠ **WARNING!**



Any product or assembly marked with these symbols indicates that the component exceeds the recommended weight for one individual to handle safely.

---

## Rack warnings and precautions

Ensure that precautions have been taken to provide for rack stability and safety. It is important to follow these precautions providing for rack stability and safety, and to protect both personnel and property. Follow all cautions and warnings included in the installation instructions.

---

⚠ **WARNING!**

To reduce the risk of personal injury or damage to the equipment:

- Observe local occupational safety requirements and guidelines for heavy equipment handling.
  - Obtain adequate assistance to lift and stabilize the product during installation or removal.
  - Extend the leveling jacks to the floor.
  - Rest the full weight of the rack on the leveling jacks.
  - Attach stabilizing feet to the rack if it is a single-rack installation.
  - Ensure the racks are coupled in multiple-rack installations.
  - Fully extend the bottom stabilizers on the equipment. Ensure that the equipment is properly supported/braced when installing options and boards.
  - Be careful when sliding rack components with slide rails into the rack. The slide rails could pinch your fingertips.
  - Ensure that the rack is adequately stabilized before extending a rack component with slide rails outside the rack. Extend only one component at a time. A rack could become unstable if more than one component is extended for any reason.
- 

---

⚠ **WARNING!**

Verify that the AC power supply branch circuit that provides power to the rack is not overloaded. Overloading AC power to the rack power supply circuit increases the risk of personal injury, fire, or damage to the equipment. The total rack load should not exceed 80 percent of the branch circuit rating. Consult the electrical authority having jurisdiction over your facility wiring and installation requirements.

---

# Device warnings and precautions

---

## WARNING!

To reduce the risk of electric shock or damage to the equipment:

- Allow the product to cool before removing covers and touching internal components.
  - Do not disable the power cord grounding plug. The grounding plug is an important safety feature.
  - Plug the power cord into a grounded (earthed) electrical outlet that is easily accessible at all times.
  - Disconnect power from the device by unplugging the power cord from either the electrical outlet or the device.
  - Do not use non-conductive tools that could bridge live parts.
  - Remove all watches, rings, or loose jewelry when working in hot-plug areas of an energized device.
  - Install the device in a controlled access location where only qualified personnel have access to the device.
  - Power off the equipment and disconnect power to all AC power cords before removing any access covers for non-hot-pluggable areas.
  - Do not replace non-hot-pluggable components while power is applied to the product. Power off the device and then disconnect all AC power cords.
  - Do not exceed the level of repair specified in the procedures in the product documentation. All troubleshooting and repair procedures are detailed to allow only subassembly or module-level repair. Because of the complexity of the individual boards and subassemblies, do not attempt to make repairs at the component level or to make modifications to any printed wiring board. Improper repairs can create a safety hazard.
- 

---

## WARNING!

To reduce the risk of personal injury or damage to the equipment, the installation of non-hot-pluggable components should be performed only by individuals who are qualified in servicing computer equipment, knowledgeable about the procedures and precautions, and trained to deal with products capable of producing hazardous energy levels.

---

---

## WARNING!

To reduce the risk of personal injury or damage to the equipment, observe local occupational health and safety requirements and guidelines for manually handling material.

---

---

## CAUTION:

Protect the installed solution from power fluctuations and temporary interruptions with a regulating Uninterruptible Power Supply (UPS). This device protects the hardware from damage caused by power surges and voltage spikes, and keeps the system in operation during a power failure.

---

---

△ **CAUTION:**

To properly ventilate the system, you must provide at least 7.6 centimeters (3.0 inches) of clearance at the front and back of the device.

---

---

△ **CAUTION:**

Schedule physical configuration changes during periods of low or no activity. If the system is performing rebuilds, RAID migrations, array expansions LUN expansions, or experiencing heavy I/O, avoid physical configuration changes such as adding or replacing hard drives or hot-plugging a controller or any other component. For example, hot-adding or replacing a controller while under heavy I/O could cause a momentary pause, performance decrease, or loss of access to the device while the new controller is starting up. When the controller completes the startup process, full functionality is restored.

---

---

△ **CAUTION:**

Before replacing a hot-pluggable component, ensure that steps have been taken to prevent loss of data.

---



---

# D Regulatory compliance and safety

## Regulatory compliance identification numbers

For the purpose of regulatory compliance certifications and identification, this product has been assigned a unique regulatory model number. The regulatory model number can be found on the product nameplate label, along with all required approval markings and information. When requesting compliance information for this product, always refer to this regulatory model number. The regulatory model number is not the marketing name or model number of the product.

## Federal Communications Commission notice

Part 15 of the Federal Communications Commission (FCC) Rules and Regulations has established Radio Frequency (RF) emission limits to provide an interference-free radio frequency spectrum. Many electronic devices, including computers, generate RF energy incidental to their intended function and are, therefore, covered by these rules. These rules place computers and related peripheral devices into two classes, A and B, depending upon their intended installation. Class A devices are those that could reasonably be expected to be installed in a business or commercial environment. Class B devices are those that could reasonably be expected to be installed in a residential environment (personal computers, for example). The FCC requires devices in both classes to bear a label indicating the interference potential of the device as well as additional operating instructions for the user.

The rating label on the device shows which class (A or B) the equipment falls into. Class B devices have an FCC logo or FCC ID on the label. Class A devices do not have an FCC logo or FCC ID on the label. Once the class of the device is determined, refer to the following corresponding statement.

### Class A equipment

This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, could cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at personal expense.

### Class B equipment

This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, could cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio or television technician for help.

## Declaration of conformity for products marked with the FCC logo, United States only

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device could not cause harmful interference, and (2) this device must accept any interference received, including interference that could cause undesired operation.

For questions regarding your product, contact:

Hewlett-Packard Company

P. O. Box 692000, Mail Stop 530113

Houston, Texas 77269-2000

Or, call

1-800- 652-6672

For questions regarding this FCC declaration, contact:

Hewlett-Packard Company

P. O. Box 692000, Mail Stop 510101

Houston, Texas 77269-2000

Or, call

(281) 514-3333

To identify this product, refer to the Part, Series, or Model number found on the product.

## Modifications

The FCC requires the user to be notified that any changes or modifications made to this device that are not expressly approved by Hewlett-Packard Company could void the user's authority to operate the equipment.

## Cables

Connections to this device must be made with shielded cables with metallic RFI/EMI connector hoods in order to maintain compliance with FCC Rules and Regulations.

## Laser compliance

This product might be provided with an optical storage device (that is, CD or DVD drive) and/or fiber optic transceiver. Each of these devices contain a laser that is classified as a Class 1 Laser Product in accordance with US FDA regulations and the IEC 60826-1. The product does not emit hazardous laser radiation.



---

⚠ **WARNING!**

Use of controls or adjustments, or performance of procedures other than those specified herein, or in the laser product's installation guide, could result in hazardous radiation exposure. To reduce the risk of exposure to hazardous radiation:

- Do not try to open the module enclosure. There are no user-serviceable components inside.
  - Do not operate controls, make adjustments, or perform procedures to the laser device, other than those specified herein.
  - Allow only HP-authorized service technicians to repair the unit.
- 

The Center for Devices and Radiological Health (CDRH) of the U.S. Food and Drug Administration implemented regulations for laser products on August 2, 1976. These regulations apply to laser products manufactured from August 1 1976. Compliance is mandatory for products marketed in the United States.

## International notices and statements

### Canadian notice (Avis Canadien)

#### Class A equipment

This Class A digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

#### Class B equipment

This Class B digital apparatus meets all requirements of the Canadian Interference-Causing Equipment Regulations.

Cet appareil numérique de la classe B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

### European Union notice

**CE** Products bearing the CE marking comply with the EMC Directive (89/336/EEC) and the Low Voltage Directive (73/23/EEC) issued by the Commission of the European Community and if this product has telecommunication functionality, the R&TTE Directive (1999/5/EC).

Compliance with these directives implies conformity to the following European Norms (in parentheses are the equivalent international standards and regulations):

- EN 55022 (CISPR 22) - Electromagnetic Interference
- EN55024 (IEC61000-4-2, 3, 4, 5, 6, 8, 11) - Electromagnetic Immunity
- EN61000-3-2 (IEC61000-3-2) - Power Line Harmonics
- EN61000-3-3 (IEC61000-3-3) - Power Line Flicker
- EN 60950 (IEC 60950) - Product Safety

## BSMI notice

**警告使用者:**

這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

## Japanese notice

[illegible]

## Korean notice (A&B)

## Class A equipment

**A급 기기 (업무용 정보통신기기)**

이 기기는 업무용으로 전자파적합등록을 한 기기이오니  
판매자 또는 사용자는 이 점을 주의하시기 바라며, 만약  
잘못판매 또는 구입하였을 때에는 가정용으로 교환하시기  
바랍니다.

## Class B equipment

**B급 기기 (가정용 정보통신기기)**

이 기기는 가정용으로 전자파적합등록을 한 기기로서  
주거지역에서는 물론 모든지역에서 사용할 수 있습니다.

# Safety


## Battery Replacement notice

---

### WARNING!


The computer contains an internal lithium manganese dioxide, a vanadium pentoxide, or an alkaline battery pack. A risk of fire and burns exists if the battery pack is not properly handled. To reduce the risk of personal injury:

- Do not attempt to recharge the battery.
  - Do not expose the battery to temperatures higher than 60°C (140°F).
  - Do not disassemble, crush, puncture, short external contacts, or dispose of in fire or water.
- 

 Batteries, battery packs, and accumulators should not be disposed of together with the general household waste. To forward them to recycling or proper disposal, please use the public collection system or return them to HP, an authorized HP Partner, or their agents.

For more information about battery replacement or proper disposal, contact an authorized reseller or an authorized service provider.

## Taiwan Battery Recycling Notice



The Taiwan EPA requires dry battery manufacturing or importing firms in accordance with Article 15 of the Waste Disposal Act to indicate the recovery marks on the batteries used in sales, giveaway or promotion. Contact a qualified Taiwanese recycler for proper battery disposal.

## Power cords

The power cord set must meet the requirements for use in the country where the product was purchased. If the product is to be used in another country, purchase a power cord that is approved for use in that country.

The power cord must be rated for the product and for the voltage and current marked on the product electrical ratings label. The voltage and current rating of the cord should be greater than the voltage and current rating marked on the product. In addition, the diameter of the wire must be a minimum of 1.00 mm<sup>2</sup> or 18 AWG, and the length of the cord must be between 1.8 m (6 ft) and 3.6 m (12 ft). If you have questions about the type of power cord to use, contact an HP-authorized service provider.

---

### NOTE:

Route power cords so that they will not be walked on and cannot be pinched by items placed upon or against them. Pay particular attention to the plug, electrical outlet, and the point where the cords exit from the product.

---

## Japanese Power Cord notice

製品には、同梱された電源コードをお使い下さい。  
同梱された電源コードは、他の製品では使用出来ません。

## Electrostatic discharge

To prevent damage to the system, be aware of the precautions you need to follow when setting up the system or handling parts. A discharge of static electricity from a finger or other conductor could damage system boards or other static-sensitive devices. This type of damage could reduce the life expectancy of the device.

### Preventing electrostatic discharge

To prevent electrostatic damage, observe the following precautions:

- Avoid hand contact by transporting and storing products in static-safe containers.
- Keep electrostatic-sensitive parts in their containers until they arrive at static-free workstations.
- Place parts on a grounded surface before removing them from their containers.
- Avoid touching pins, leads, or circuitry.
- Always be properly grounded when touching a static-sensitive component or assembly.

### Grounding methods

There are several methods for grounding. Use one or more of the following methods when handling or installing electrostatic-sensitive parts:

- Use a wrist strap connected by a ground cord to a grounded workstation or computer chassis. Wrist straps are flexible straps with a minimum of 1 megohm  $\pm$  10 percent resistance in the ground cords. To provide proper grounding, wear the strap snug against the skin.
- Use heel straps, toe straps, or boot straps at standing workstations. Wear the straps on both feet when standing on conductive floors or dissipating floor mats.
- Use conductive field service tools.
- Use a portable field service kit with a folding static-dissipating work mat.

If you do not have any of the suggested equipment for proper grounding, have an authorized reseller install the part.



#### NOTE:

For more information on static electricity, or for assistance with product installation, contact your authorized reseller.

---

# Waste Electrical and Electronic Equipment directive

## Czechoslovakian notice

### Likvidace zařízení soukromými domácími uživateli v Evropské unii



■ Tento symbol na produktu nebo balení označuje výrobek, který nesmí být vyhozen spolu s ostatním domácím odpadem. Povinností uživatele je předat takto označený odpad na předem určené sběrné místo pro recyklaci elektrických a elektronických zařízení. Okamžité třídění a recyklace odpadu pomůže uchovat přírodní prostředí a zajistí takový způsob recyklace, který ochrání zdraví a životní prostředí člověka. Další informace o možnostech odevzdání odpadu k recyklaci získáte na příslušném obecním nebo městském úřadě, od firmy zabývající se sběrem a svozem odpadu nebo v obchodě, kde jste produkt zakoupili.

## Danish notice

### Bortskaffelse af affald fra husstande i den Europæiske Union



■ Hvis produktet eller dets emballage er forsynet med dette symbol, angiver det, at produktet ikke må bortskaffes med andet almindeligt husholdningsaffald. I stedet er det dit ansvar at bortskaffe kasseret udstyr ved at aflevere det på den kommunale genbrugsstation, der forestår genvinding af kasseret elektrisk og elektronisk udstyr. Den centrale modtagelse og genvinding af kasseret udstyr i forbindelse med bortskaffelsen bidrager til bevarelse af naturlige ressourcer og sikrer, at udstyret genvindes på en måde, der beskytter både mennesker og miljø. Yderligere oplysninger om, hvor du kan aflevere kasseret udstyr til genvinding, kan du få hos kommunen, den lokale genbrugsstation eller i den butik, hvor du købte produktet.

## Dutch notice

### Verwijdering van afgedankte apparatuur door privé-gebruikers in de Europese Unie



■ Dit symbool op het product of de verpakking geeft aan dat dit product niet mag worden gedeponeerd bij het normale huishoudelijke afval. U bent zelf verantwoordelijk voor het inleveren van uw afgedankte apparatuur bij een inzamelingspunt voor het recyclen van oude elektrische en elektronische apparatuur. Door uw oude apparatuur apart aan te bieden en te recyclen, kunnen natuurlijke bronnen worden behouden en kan het materiaal worden hergebruikt op een manier waarmee de volksgezondheid en het milieu worden beschermd. Neem contact op met uw gemeente, het afvalinzamelingsbedrijf of de winkel waar u het product hebt gekocht voor meer informatie over inzamelingspunten waar u oude apparatuur kunt aanbieden voor recycling.

## English notice

### Disposal of waste equipment by users in private household in the European Union



■ This symbol on the product or on its packaging indicates that this product must not be disposed of with your other household waste. Instead, it is your responsibility to dispose of your waste equipment by handing it over to a designated collection point for recycling of waste electrical and electronic equipment. The separate collection and recycling of your waste equipment at the time of disposal will help to conserve natural resources and ensure that it is recycled in a manner that protects human health and the environment. For more information about where you can drop off your waste equipment for recycling, please contact your local city office, your household waste disposal service, or the shop where you purchased the product.

## Estonian notice

### Seadmete jäätmete kõrvaldamine eramajapidamistes Euroopa Liidus



■ See tootel või selle pakendil olev sümbol näitab, et kõnealust toodet ei tohi koos teiste majapidamisjätmetega kõrvaldada. Teie kohus on oma seadmete jäätmed kõrvaldada, viies need elektri- ja elektroonikaseadmete jäätmete ringlussevõtmiseks selleks ettenähtud kogumispunkti. Seadmete jäätmete eraldi kogumine ja ringlussevõtmine kõrvaldamise ajal aitab kaitsta loodusvarasid ning tagada, et ringlussevõtmine toimub viisil, mis kaitseb inimeste tervist ning keskkonda. Lisateabe saamiseks selle kohta, kuhu oma seadmete jäätmed ringlussevõtmiseks viia, võtke palun ühendust oma kohaliku linnakantselei, majapidamisjätmete kõrvaldamise teenistuse või kauplusega, kust Te toote ostsite.

## Finnish notice

### Laitteiden hävittäminen kotitalouksissa Euroopan unionin alueella



■ Jos tuotteessa tai sen pakkauksessa on tämä merkki, tuotetta ei saa hävittää kotitalousjätteiden mukana. Tällöin hävitettävä laite on toimitettava sähkölaitteiden ja elektronisten laitteiden kierrätyspisteeseen. Hävitettävien laitteiden erillinen käsittely ja kierrätys auttavat säästämään luonnonvaroja ja varmistamaan, että laite kierrätetään tavalla, joka estää terveyshaitat ja suojelee luontoa. Lisätietoja paikoista, joihin hävitettävät laitteet voi toimittaa kierrätettäväksi, saa ottamalla yhteyttä jätehuoltoon tai liikkeeseen, josta tuote on ostettu.

## French notice

### Élimination des appareils mis au rebut par les ménages dans l'Union européenne



■ Le symbole apposé sur ce produit ou sur son emballage indique que ce produit ne doit pas être jeté avec les déchets ménagers ordinaires. Il est de votre responsabilité de mettre au rebut vos appareils en les déposant dans les centres de collecte publique désignés pour le recyclage des équipements électriques et électroniques. La collecte et le recyclage de vos appareils mis au rebut indépendamment du reste des déchets contribue à la préservation des ressources naturelles et garantit que ces appareils seront recyclés dans le respect de la santé humaine et de l'environnement. Pour obtenir plus d'informations sur les centres de collecte et de recyclage des appareils mis au rebut, veuillez contacter les autorités locales de votre région, les services de collecte des ordures ménagères ou le magasin dans lequel vous avez acheté ce produit.

## German notice

### Entsorgung von Altgeräten aus privaten Haushalten in der EU



■ Das Symbol auf dem Produkt oder seiner Verpackung weist darauf hin, dass das Produkt nicht über den normalen Hausmüll entsorgt werden darf. Benutzer sind verpflichtet, die Altgeräte an einer Rücknahmestelle für Elektro- und Elektronik-Altgeräte abzugeben. Die getrennte Sammlung und ordnungsgemäße Entsorgung Ihrer Altgeräte trägt zur Erhaltung der natürlichen Ressourcen bei und garantiert eine Wiederverwertung, die die Gesundheit des Menschen und die Umwelt schützt. Informationen dazu, wo Sie Rücknahmestellen für Ihre Altgeräte finden, erhalten Sie bei Ihrer Stadtverwaltung, den örtlichen Müllentsorgungsbetrieben oder im Geschäft, in dem Sie das Gerät erworben haben.

## Greek notice

### Απορρίψη άχρηστου εξοπλισμού από χρήστες σε ιδιωτικά νοικοκυριά στην Ευρωπαϊκή Ένωση



■ Το σύμβολο αυτό στο προϊόν ή τη συσκευασία του υποδεικνύει ότι το συγκεκριμένο προϊόν δεν πρέπει να διατίθεται μαζί με τα άλλα οικιακά σας απορρίμματα. Αντίθετα, είναι δική σας ευθύνη να απορρίψετε τον άχρηστο εξοπλισμό σας παραδίδοντάς τον σε καθορισμένο σημείο συλλογής για την ανακύκλωση άχρηστου ηλεκτρικού και ηλεκτρονικού εξοπλισμού. Η ξεχωριστή συλλογή και ανακύκλωση του άχρηστου εξοπλισμού σας κατά την απορρίψη θα συμβάλει στη διατήρηση των φυσικών πόρων και θα διασφαλίσει ότι η ανακύκλωση γίνεται με τρόπο που προστατεύει την ανθρώπινη υγεία και το περιβάλλον. Για περισσότερες πληροφορίες σχετικά με το πού μπορείτε να παραδώσετε τον άχρηστο εξοπλισμό σας για ανακύκλωση, επικοινωνήστε με το αρμόδιο τοπικό γραφείο, την τοπική υπηρεσία διάθεσης οικιακών απορριμμάτων ή το κατάστημα όπου αγοράσατε το προϊόν.



## Hungarian notice

### Készülékek magánháztartásban történő selejtezése az Európai Unió területén



■ A készüléken, illetve a készülék csomagolásán látható azonos szimbólum annak jelzésére szolgál, hogy a készülék a selejtezés során az egyéb háztartási hulladéktól eltérő módon kezelendő. A vásárló a hulladékká vált készüléket köteles a kijelölt gyűjtőhelyre szállítani az elektromos és elektronikai készülékek újrahasznosítása céljából. A hulladékká vált készülékek selejtezés kori begyűjtése és újrahasznosítása hozzájárul a természeti erőforrások megőrzéséhez, valamint biztosítja a selejtezett termékek környezetre és emberi egészségre nézve biztonságos feldolgozását. A begyűjtés pontos helyéről bővebb tájékoztatást a lakhelye szerint illetékes önkormányzattól, az illetékes személtakarító vállalatától, illetve a terméket elárúsító helyen kaphat.

## Italian notice

### Smaltimento delle apparecchiature da parte di privati nel territorio dell'Unione Europea



■ Questo simbolo presente sul prodotto o sulla sua confezione indica che il prodotto non può essere smaltito insieme ai rifiuti domestici. È responsabilità dell'utente smaltire le apparecchiature consegnandole presso un punto di raccolta designato al riciclo e allo smaltimento di apparecchiature elettriche ed elettroniche. La raccolta differenziata e il corretto riciclo delle apparecchiature da smaltire permette di proteggere la salute degli individui e l'ecosistema. Per ulteriori informazioni relative ai punti di raccolta delle apparecchiature, contattare l'ente locale per lo smaltimento dei rifiuti, oppure il negozio presso il quale è stato acquistato il prodotto.

## Latvian notice

### Nolietotu iekārtu iznīcināšanas noteikumi lietotājiem Eiropas Savienības privātajās mājāsaimniecībās



■ Šāds simbols uz izstrādājuma vai uz tā iesaiņojuma norāda, ka šo izstrādājumu nedrīkst izmest kopā ar citiem sadzīves atkritumiem. Jūs atbildat par to, lai nolietotās iekārtas tiktu nodotas speciāli iekārtotos punktos, kas paredzēti izmantoto elektrisko un elektronisko iekārtu savākšanai otrreizējai pārstrādei. Atsevišķa nolietoto iekārtu savākšana un otrreizējā pārstrāde palīdzēs saglabāt dabas resursus un garantēs, ka šīs iekārtas tiks otrreizēji pārstrādātas tādā veidā, lai pasargātu vidi un cilvēku veselību. Lai uzzinātu, kur nolietotās iekārtas var izmest otrreizējai pārstrādei, jāvērsas savas dzīves vietas pašvaldībā, sadzīves atkritumu savākšanas dienestā vai veikalā, kurā izstrādājums tika nopirkts.



## Lithuanian notice

### Vartotojų iš privačių namų ūkių įrangos atliekų šalinimas Europos Sąjungoje



■ Šis simbolis ant gaminio arba jo pakuotės rodo, kad šio gaminio šalinti kartu su kitomis namų ūkio atliekomis negalima. Šalintinas įrangos atliekas privalote pristatyti į specialią surinkimo vietą elektros ir elektroninės įrangos atliekoms perdirbti. Atskirai surenkamos ir perdirbamos šalintinos įrangos atliekos padės saugoti gamtinius išteklius ir užtikrinti, kad jos bus perdirbtos tokiu būdu, kuris nekenkia žmonių sveikatai ir aplinkai. Jeigu norite sužinoti daugiau apie tai, kur galima pristatyti perdirbtinas įrangos atliekas, kreipkitės į savo seniūniją, namų ūkio atliekų šalinimo tarnybą arba parduotuvę, kurioje įsigijote gaminį.

## Polish notice

### Pozbywanie się zużytego sprzętu przez użytkowników w prywatnych gospodarstwach domowych w Unii Europejskiej



■ Ten symbol na produkcie lub jego opakowaniu oznacza, że produktu nie wolno wyrzucać do zwykłych pojemników na śmieci. Obowiązkiem użytkownika jest przekazanie zużytego sprzętu do wyznaczonego punktu zbiórki w celu recyklingu odpadów powstałych ze sprzętu elektrycznego i elektronicznego. Osobna zbiórka oraz recykling zużytego sprzętu pomogą w ochronie zasobów naturalnych i zapewnią ponowne wprowadzenie go do obiegu w sposób chroniący zdrowie człowieka i środowisko. Aby uzyskać więcej informacji o tym, gdzie można przekazać zużyty sprzęt do recyklingu, należy się skontaktować z urzędem miasta, zakładem gospodarki odpadami lub sklepem, w którym zakupiono produkt.

## Portuguese notice

### Descarte de Lixo Elétrico na Comunidade Européia



■ Este símbolo encontrado no produto ou na embalagem indica que o produto não deve ser descartado no lixo doméstico comum. É responsabilidade do cliente descartar o material usado (lixo elétrico), encaminhando-o para um ponto de coleta para reciclagem. A coleta e a reciclagem seletivas desse tipo de lixo ajudarão a conservar as reservas naturais; sendo assim, a reciclagem será feita de uma forma segura, protegendo o ambiente e a saúde das pessoas. Para obter mais informações sobre locais que reciclam esse tipo de material, entre em contato com o escritório da HP em sua cidade, com o serviço de coleta de lixo ou com a loja em que o produto foi adquirido.

## Slovakian notice

### Likvidácia vyradených zariadení v domácnostiach v Európskej únii



■ Symbol na výrobku alebo jeho balení označuje, že daný výrobok sa nesmie likvidovať s domovým odpadom. Povinnosťou spotrebiteľa je odovzdať vyradené zariadenie v zbernom mieste, ktoré je určené na recykláciu vyradených elektrických a elektronických zariadení. Separovaný zber a recyklácia vyradených zariadení prispieva k ochrane prírodných zdrojov a zabezpečuje, že recyklácia sa vykonáva spôsobom chrániacim ľudské zdravie a životné prostredie. Informácie o zberných miestach na recykláciu vyradených zariadení vám poskytne miestne zastupiteľstvo, spoločnosť zabezpečujúca odvoz domového odpadu alebo obchod, v ktorom ste si výrobok zakúpili.

## Slovenian notice

### Odstranjevanje odslužene opreme uporabnikov v zasebnih gospodinjstvih v Evropski uniji



■ Ta znak na izdelku ali njegovi embalaži pomeni, da izdelka ne smete odvreči med gospodinjske odpadke. Nasprotno, odsluženo opremo morate predati na zbirališče, pooblaščen za recikliranje odslužene električne in elektronske opreme. Ločeno zbiranje in recikliranje odslužene opreme prispeva k ohranjanju naravnih virov in zagotavlja recikliranje te opreme na zdravju in okolju neškodljiv način. Za podrobnejše informacije o tem, kam lahko odpeljete odsluženo opremo na recikliranje, se obrnite na pristojni organ, komunalno službo ali trgovino, kjer ste izdelek kupili.

## Spanish notice

### Eliminación de residuos de equipos eléctricos y electrónicos por parte de usuarios particulares en la Unión Europea



■ Este símbolo en el producto o en su envase indica que no debe eliminarse junto con los desperdicios generales de la casa. Es responsabilidad del usuario eliminar los residuos de este tipo depositándolos en un "punto limpio" para el reciclado de residuos eléctricos y electrónicos. La recogida y el reciclado selectivos de los residuos de aparatos eléctricos en el momento de su eliminación contribuirá a conservar los recursos naturales y a garantizar el reciclado de estos residuos de forma que se proteja el medio ambiente y la salud. Para obtener más información sobre los puntos de recogida de residuos eléctricos y electrónicos para reciclado, póngase en contacto con su ayuntamiento, con el servicio de eliminación de residuos domésticos o con el establecimiento en el que adquirió el producto.

## Swedish notice

### Bortskaffande av avfallsprodukter från användare i privathushåll inom Europeiska Unionen



Om den här symbolen visas på produkten eller förpackningen betyder det att produkten inte får slängas på samma ställe som hushållssopor. I stället är det ditt ansvar att bortskaffa avfallet genom att överlämna det till ett uppsamlingsställe avsett för återvinning av avfall från elektriska och elektroniska produkter. Separat insamling och återvinning av avfallet hjälper till att spara på våra naturresurser och gör att avfallet återvinns på ett sätt som skyddar människors hälsa och miljön. Kontakta ditt lokala kommunkontor, din närmsta återvinningsstation för hushållsavfall eller affären där du köpte produkten för att få mer information om var du kan lämna ditt avfall för återvinning.



---

# Glossary

<b>ACE</b>	Access control entry.
<b>ACL</b>	Access control list.
<b>ADS</b>	Active Directory Service.
<b>ALB</b>	Advanced load balancing.
<b>BMC</b>	Baseboard Management Configuration.
<b>CIFS</b>	Common Internet File System. The protocol used in Windows environments for shared folders.
<b>CLI</b>	Command-line interface. An interface comprised of various commands which are used to control operating system responses.
<b>CSR</b>	Customer self repair.
<b>DAS</b>	Direct attach storage. A dedicated storage device that connects directly to one or more servers.
<b>DNS</b>	Domain name system.
<b>FTP</b>	File Transfer Protocol.
<b>GSI</b>	Global service indicator.
<b>HA</b>	High availability.
<b>HBA</b>	host bus adapter.
<b>HCA</b>	Host channel adapter.
<b>HDD</b>	Hard disk drive.
<b>IAD</b>	HP X9000 Software Administrative Daemon.
<b>iLO</b>	Integrated Lights-Out.
<b>IML</b>	Initial microcode load.
<b>IOPS</b>	I/Os per second.
<b>IPMI</b>	Intelligent Platform Management Interface.
<b>JBOD</b>	Just a bunch of disks.
<b>KVM</b>	Keyboard, video, and mouse.
<b>LUN</b>	Logical unit number.

<b>MTU</b>	Maximum Transmission Unit.
<b>NAS</b>	Network attached storage.
<b>NFS</b>	Network file system. The protocol used in most UNIX environments to share folders or mounts.
<b>NIC</b>	Network interface card. A device that handles communication between a device and other devices on a network.
<b>NTP</b>	Network Time Protocol. A protocol that enables the storage system's time and date to be obtained from a network-attached server, keeping multiple hosts and storage devices synchronized.
<b>OA</b>	HP Onboard Administrator.
<b>OFED</b>	OpenFabrics Enterprise Distribution.
<b>OSD</b>	On-screen display.
<b>OU</b>	Active Directory Organizational Units.
<b>RO</b>	Read-only access.
<b>RPC</b>	Remote Procedure Call.
<b>RW</b>	Read-write access.
<b>SAN</b>	Storage area network. A network of storage devices available to one or more servers.
<b>SAS</b>	Serial Attached SCSI.
<b>SELinux</b>	Security-Enhanced Linux.
<b>SFU</b>	Microsoft Services for UNIX.
<b>SID</b>	Secondary controller identifier number.
<b>SNMP</b>	Simple Network Management Protocol.
<b>TCP/IP</b>	Transmission Control Protocol/Internet Protocol.
<b>UDP</b>	User Datagram Protocol.
<b>UFM</b>	Voltaire's Unified Fabric Manager client software.
<b>UID</b>	Unit identification.
<b>USM</b>	SNMP User Security Model.
<b>VACM</b>	SNMP View Access Control Model.
<b>VC</b>	HP Virtual Connect.
<b>VIF</b>	Virtual interface.
<b>WINS</b>	Windows Internet Naming Service.
<b>WWN</b>	World Wide Name. A unique identifier assigned to a Fibre Channel device.

<b>WWNN</b>	World wide node name. A globally unique 64-bit identifier assigned to each Fibre Channel node process.
<b>WWPN</b>	World wide port name. A unique 64-bit address used in a FC storage network to identify each device in a FC network.





---

# Index

## A

- agile management console, 25
- AutoPass, 91

## B

- backups
  - file systems, 43
  - management console configuration, 43
  - NDMP applications, 43

## C

- Class A equipment, 143
- Class B equipment, 143
- cluster
  - events, monitor, 51
  - health checks, 52
  - license key, 91
  - license, view, 91
  - log files, 55
  - operating statistics, 56
  - version numbers, view, 97
- cluster interface
  - change IP address, 66
  - change network, 66
  - defined, 63
- components
  - returning, 101
- contacting HP, 123
- customer self repair, 124

## D

- declaration of conformity, 144
- document
  - related information, 123

## E

- email event notification, 37

- events, cluster

- add SNMPv3 users and groups, 41
  - configure email notification, 37
  - configure SNMP agent, 39
  - configure SNMP notification, 38
  - configure SNMP trapsinks, 40
  - define MIB views, 41
  - delete SNMP configuration elements, 42
  - enable or disable email notification, 38
  - list email notification settings, 38
  - list SNMP configuration, 42
  - monitor, 51
  - remove, 52
  - view, 52

## F

- FCC logo, 144
- file serving node
  - recover, 107
- file serving nodes
  - configure power sources for failover, 28
  - dissociate power sources, 29
  - fail back, 30
  - fail over manually, 30
  - health checks, 52
  - identify standbys, 27
  - maintain consistency with configuration database, 99
  - migrate segments, 60
  - monitor status, 51
  - operational states, 51
  - power management, 58
  - prefer a user network interface, 65
  - run health check, 100
  - start or stop processes, 59
  - troubleshooting, 98
  - tune, 59
  - view process status, 59
- file systems
  - segments
    - migrate, 60

## G

grounding  
methods, 137

## H

hazardous conditions  
symbols on equipment, 137

### HBA

delete HBAs, 34  
delete standby port pairings, 34  
discover, 33  
identify standby-paired ports, 34  
list information, 34  
monitor for high availability, 33  
monitoring, turn on or off, 34

health check reports, 53

### help

obtaining, 123

### High Availability

agile management console, 25  
automated failover, turn on or off, 30  
check configuration, 35  
defined, 26  
delete network interface monitors, 32  
delete network interface standbys, 33  
delete power sources, 29  
detailed configuration report, 36  
dissociate power sources, 29  
fail back a node, 30  
failover a node manually, 30  
failover protection, 10  
HBA monitoring, turn on or off, 34  
identify network interface monitors, 32  
identify network interface standbys, 32  
identify standby-paired HBA ports, 34  
identify standbys for file serving nodes, 27  
power management for nodes, 58  
set up automated failover, 27  
set up HBA monitor, 33  
set up manual failover, 30  
set up network interface monitoring, 31  
set up power sources, 28  
summary configuration report, 35  
troubleshooting, 98

### hostgroups, 47

add domain rule, 48  
add X9000 client, 48  
create hostgroup tree, 48  
delete, 49  
prefer a user network interface, 65  
view, 49

### HP

technical support, 123

## I

international notices and statements, 145

### IP address

change for cluster interface, 66  
change for X9000 client, 66

## L

labels, symbols on equipment, 137

laser compliance, 144

Linux X9000 clients, upgrade, 88

loading rack, warning, 138

## M

### management console

agile, 25  
back up configuration, 43  
convert to agile configuration, 69  
failover, 25  
X9000 client access, 18  
management console CLI, 18  
management console GUI  
change password, 19  
customize, 17  
Details page, 16  
Navigator, 16  
open, 14  
view events, 52

## N

### NDMP backups, 43

cancel sessions, 45  
configure NDMP parameters, 44  
rescan for new devices, 46  
start or stop NDMP Server, 45  
view events, 46  
view sessions, 45  
view tape and media changer devices, 46

### network interfaces

add routing table entries, 67  
bonded and virtual interfaces, 63  
defined, 63  
delete, 67  
delete monitors, 32  
delete routing table entries, 67  
delete standbys, 33  
identify monitors, 32  
identify standbys, 32  
set up monitoring, 31  
viewing, 67

## P

passwords, change  
GUI password, 19

## Q

QuickRestoreDVD, 107

## R

rack stability  
warning, 124  
regulatory compliance, 143  
related documentation, 123  
returning components, 101  
routing table entries  
add, 67  
delete, 67

## S

segments  
evacuate from cluster, 61  
migrate, 60  
shared ssh keys, configure, 97  
SNMP event notification, 38  
SNMP MIB, 41  
storage, remove from cluster, 61  
support tickets, 95  
symbols  
on equipment, 137  
System Configuration Wizard, 108  
system recovery, 107  
System Configuration Wizard, 108

## T

technical support  
HP, 123  
service locator website, 124

## U

upgrades  
firmware, 93  
Linux X9000 clients, 88  
Windows X9000 clients, 89  
X9000 Software, 75

user network interface  
add, 63  
configuration rules, 66  
defined, 63  
identify for X9000 clients, 64  
modify, 64  
prefer, 65  
unprefer, 65

## W

warning  
rack stability, 124  
warnings  
loading rack, 138  
Waste Electrical and Electronic Equipment  
directive, 149  
websites  
customer self repair, 124  
HP, 124  
Windows X9000 clients, upgrade, 89

## X

X9000 clients  
add to hostgroup, 48  
change IP address, 66  
identify a user network interface, 64  
interface to management console, 18  
migrate segments, 60  
monitor status, 51  
prefer a user network interface, 65  
start or stop processes, 59  
troubleshooting, 98  
tune, 59  
tune locally, 60  
view process status, 59  
X9000 Software  
start, 58  
X9000 Software  
shut down, 57  
X9000 Software upgrade, 75  
agile offline upgrade, 85  
agile online upgrade, 81  
agile upgrade, defined, 76  
standard offline upgrade, 79  
standard online upgrade, 77  
standard upgrade, defined, 76